Decreasing Marginal Rate of Growth and its Business Cycle Consequences

Moshe Elitzur[a,*], Oskar Zorrilla[b], Željko Ivezić[c], David Zilberman[d]

[a]*Dept of Physics & Astronomy, Univ of Kentucky, Lexington, KY*
[b]*Economics Department, US Naval Academy, Annapolis, Md*
[c]*Dept of Astronomy, University of Washington, Seattle, WA*
[d]*Dept of Agricultural & Resource Economics, Univ of California, Berkeley, CA*

Abstract

We introduce a new detrending method and employ it to decompose 234 years of US GDP into a log concave growth trend and stationary fluctuations. The mean magnitude of model deviation from data is 2.75%. Initially exponential, the trend rises toward a linear asymptote where 1.05 trillion 2020 USD are added every year to the GDP. The fluctuations around the trend are directly coupled for 2 years; their autocorrelations exceed 0.5 for 10 years, remaining above the noise level for up to 28 years. They contain white noise whose probability distribution has a Gaussian core and non-Gaussian tails. It is highly leptokurtic with a strong negative skew, reflecting an overabundance of negative outliers, i.e., disasters. The outliers are well matched with the historical record and show that US economic volatility subsided significantly after 1946. The unit root hypothesis is rejected decisively.

Keywords: growth, hindered growth, time series analysis, stochastic variability, diminishing returns

1. Introduction

US GDP grew by a factor of 4,500 from $5.24 billion in 1790, measured in 2020 USD, to $23.6 trillion in 2023 (Williamson, 2024). In addition to this overall growth, it also displays stochastic time variability, undergoing a sequence of changes that is recurrent but not periodic (Burns and Mitchell, 1946). Macroeconomics is organized around these two empirical facts. The study of economic

*Corresponding author; email: moshe@g.uky.edu

growth focuses on the trend, while the study of business cycles focuses on the stochastic variability. Decomposing GDP into its trend and cyclical components is, therefore, of prime importance. Such decomposition would be straightforward if the growth trend could be reliably described with an exponential function. However, the GDP annual growth rate shows a clear downward trend: while its 50-years mean is 4.15% for the period 1790–1840, it is only 2.63% for 1973–2023. Without an adequate description for the trend, a meaningful decomposition to distinguish growth from the business cycle is impossible; as noted by Blanchard and Fischer (1989; p 12), growth is, of course, a nonstationary process and "there are infinitely many ways of decomposing a nonstationary process as the sum of a nonstationary process (the trend) and a stationary process (the cycle)."

In this paper we propose and apply a new decomposition method. Instead of assuming a constant growth rate we allow it to be weakly decreasing while maintaining the requirement that it be well-defined everywhere. We prove mathematically that these broad requirements suffice to define an ensemble of log concave *hindering functions*, recently employed in Elitzur et al. (2020), Elitzur et al. (2021) and Elitzur (2022), and test the hypothesis that the US GDP growth trend follows one of those. To that end, our model is comprised of a deterministic growth trend described by a hindering function and a stochastic stationary component generated in an ARMA process. Searching the ensemble of hindering functions for the one providing the best fit, our approach relies on the data to identify the trend rather than impose a trend on the data.

Our approach is efficient: although the space of hindering functions is infinite, our search method yields a solution in seconds. We find that US GDP has grown during 1790–2023 at a marginally diminishing rate, tending toward asymptotic growth that is linear rather than exponential. Our approach is also parsimonious: the best-fitting trend adds *a single free parameter* to the two of exponential growth. Because the exponential belongs to the hindering functions ensemble, we are able to reject overwhelmingly the null of exponential growth with a straightforward F-test. Our approach is robust: using data only up to 2003 yields a model statistically identical to the one for the full dataset. The model derived from the partial data predicts the subsequent 2004–2023 GDP with high accuracy; in particular, it describes properly the decline of the underlying growth rate for this period, which includes the Great Recession and COVID-19 pandemic. In contrast, fitting the pre-2004 data with an exponential yields exceedingly poor predictions for the subsequent 20 years; after 2013, every data point is outside the 95% confidence interval of this model.

We find that the fluctuations, the stochastic component, are directly coupled for two years, strongly correlated when separated by up to 10 years and more weakly for up to 28 years, blending with the noise at longer separations. The ARMA residuals (shocks) are compatible with white noise. Their probability distribution has a Gaussian core and prominent non-Gaussian tails. It is highly leptokurtic with a strong negative skew, implying that extreme negative outliers, colloquially termed "Black Swans", outnumber positive ones. We find that US economic volatility has greatly subsided after World War II; every post-1946 shock is in the Gaussian core, within 2σ of the distribution center.

There is ample literature on detrending macroeconomic data. Like Razzak (1997), we too find that the Hodrick and Prescott (1981) decomposition is a smoothing, not a detrending technique. Groth et al. (2010) propose their own parametric description for log concavity. For each of their proposed functions the growth process blows up when extended back in time. This same problem afflicts other putative expressions of log concavity, including the widely used log quadratic approximation (e.g. Mendoza, 1991). Recognizing this problem, Groth et al. (2010) suggest that their proposal should only be considered a local description of the trend, applicable only for a limited period. Since our hindering functions are singularity-free everywhere, they can be considered a *global* estimation of the trend.

In their seminal work, Nelson and Plosser (1982) show that under the null of exponential growth, US output follows a random walk. Perron (1989) introduces a second class of shocks—structural breaks—that eliminate the random walk in output fluctuations. The Perron method transforms detrending into a series of local problems by breaking the data into segments and determining for each one a growth rate smaller than its precursor, creating a log concave piecewise exponential trend. Eo and Morley (2022) suggest that the Great Recession necessitates a new structural break to account for a new step-down in growth rate. As noted above, our global detrending approach predicts successfully the growth rate during 2004–2023 from the pre-2004 data with a single logarithmically concave smooth hindering function. Growth slowdown after the Great Recession, as well as before, is an organic feature of the model; it requires no special handling or any additional parameters.

Comin and Gertler (2006) suggest that in addition to a high-frequency component in accordance with Burns & Mitchell, business cycles also have a medium-frequency component that extends to longer cycles. This suggestion is aligned with the duration of ~28 years we find for stochastic temporal structure detectable above the noise level, and for the rough division at ~10 years by its

autocorrelation magnitude. The prominent non-Gaussian tails we find for the shock distribution fit into the larger literature that reconsiders the ubiquitous assumption of normality in macroeconomics. Using a panel of 24 countries over 100 years, Nakamura et al. (2013) find, like us, that the probability of large negative outliers—what the literature terms disasters—is significantly higher than that implied by a normal distribution. Gabaix (2011) shows that when firm sizes are not normally distributed, firm-level shocks can generate aggregate fluctuations.

Most importantly, our results are well matched with the historical record, which is not the case for other trends. For example, our analysis identifies successive large negative shocks for the Great Depression years, culminating in a -5.5σ shock in 1932. In contrast, the HP filter identifies no negative shocks during 1930–1933, instead finding a large *positive* one in 1931, the year the growth rate hit its all-time bottom of -9.9%. And according to the exponential trend, 70 years ago the GDP output gap turned negative and started a still-ongoing downward slide. The negative gap has been widening continuously to the point that it surpassed the Great Depression level 30 years ago, implying that the US economy has been mired for years in its longest, worst depression ever.

The paper is structured as follows. In §2 we formally define hindering functions, derive their general solution and characterize some of their properties. This places the mathematical foundations of the hindering formalism on more solid footing than previously (Elitzur et al., 2020; Elitzur, 2022). Section 3 presents detailed analysis of the US GDP evolution from 1790 to 2023. The analysis is conducted in terms of both hindering detrending (§3.1) and the HP decomposition technique (§3.2). Forecasting is discussed in §4, while §5 presents detailed analysis of the fluctuations around the trend; stochastic variability (the business cycle) is studied in §5.1, the residuals of its ARMA fit (shocks) in §5.2. We conclude by discussing our results in section §6 .

2. Hindered Growth Processes

Consider a dynamic process $Y(t) > 0$ comprised of a deterministic growth trend $Q(t)$ and a stochastic component $X(t)$. The trend has a continuous, smooth growth rate $g(t)$ that can be expressed as the generic equation of growth

$$\frac{d \ln Q}{dt} = g, \tag{2.1}$$

a differential equation for the time variation of Q. Given $g(t)$, $Q(t)$ is determined from straightforward integration that can start from arbitrarily small $Q > 0$.

Definition 1. The variable Q describes a *hindered growth process* if its growth rate g satisfies the following three conditions:

1. Monotonicity: $g \geq 0$
2. Stability: $\dot{g} \leq 0$
3. Boundedness: $\lim_{Q \to 0} g = \gamma \in \mathbb{R}_+$

where $\dot{g}$ denotes the derivative of g with respect to time.

Monotonicity is self-explanatory; a nondecreasing trend implies that any contractions of Y are due to the stochastic component, X. A monotonically-varied Q is single-valued, so g can also be considered a function of Q instead of the independent variable t. *Stability* requires that the growth rate is either constant or monotonically decreasing, implying that Q is log concave. As we show in Appendix A.1. accelerated growth, $\dot{g} > 0$, is unstable. *Boundedness* is required for a smooth, well-behaving trend function, so g must be free of singularities in the entire upper half of the (t, Q) plane and must remain finite when $Q \to 0$, or, equivalently, when $t \to -\infty$. Note that exponential growth satisfies all three conditions.

To classify growth patterns, the equation of growth (2.1) has to be transformed into a dimensionless mathematical equation by extracting intrinsic scales of Q and t. The inverse of γ defines such a time scale. Time measured in units of $1/\gamma$, namely γt, is the intrinsic independent variable of the growth process.[1] Our next step is to center the intrinsic time scale. Definition 1 restricts the growth rate to the compact range $[0, \gamma]$. A natural center occurs at the half-life of the growth rate, $g = \gamma/2$. In the segment $[\frac{1}{2}\gamma, \gamma]$, $g \simeq \gamma$ to leading order, whereas in the segment $[0, \frac{1}{2}\gamma]$, g is closer to 0 than it is to γ. We denote the time of the half-life of g by τ and the corresponding value of Q by Q_τ:

$$g(\tau) = \tfrac{1}{2}\gamma, \quad Q_\tau = Q(\tau). \tag{2.2}$$

Boundedness enables us to transform the time domain to

$$w = \gamma(t - \tau) \tag{2.3}$$

and we can then write the hindered growth process as

$$Q(w) = Q_\tau \, h(w) \tag{2.4}$$

[1] This has a direct analogue to periodic processes. The intrinsic independent variable of a periodic process with frequency ω is ωt.

where $h(w)$ is a profile function normalized to $h(0) = 1$ that describes the shape of the growth trajectory.

The growth rate g can be similarly expressed in terms of a mathematical profile function p, defined through

$$g(Q) = \frac{\gamma}{1 + p(q)} \tag{2.5}$$

with $q = Q/Q_\tau$. Note that $q(t) = h(w)$ — these functions describe the same variation in terms of different time variables. From Definition 1, p is monotonically increasing with $p(0) = 0$;[2] the definition in eq. (2.2) implies $p(1) = 1$. Substituting the profiles h for Q and p for g, the equation of growth (2.1) becomes

$$\frac{d \ln h}{dw} = \frac{1}{1 + p}; \tag{2.6}$$

note that $h' = \frac{1}{2}$ at the point $w = 0$ ($q = 1$), with the prime denoting derivative with respect to w. With the boundary condition $h(0) = 1$, the formal solution

$$\ln h(w) + \int_1^{h(w)} p(q) \frac{dq}{q} = w \tag{2.7}$$

determines the shape of the growth trajectory. The profile function p describes the deviations from purely exponential growth.

2.1. Analytic p-profiles

When the profile p is analytic (infinitely differentiable), it can be expanded in a power series

$$p(q) = \sum_{k=1}^{\infty} \lambda_k q^k \tag{2.8}$$

with $\lambda_0 = 0$ because $p(0) = 0$. A monotonically increasing p implies $\lambda_k \geq 0$ for all k, while $p(1) = 1$ (eq. 2.2) implies

$$\sum_{k=1}^{\infty} \lambda_k = 1. \tag{2.9}$$

[2] While $p \geq 0$ describes hindered growth, $-1 < p < 0$ corresponds to accelerated growth and $p < -1$ describes a contracting Q.

Theorem 1. *Every hindered growth process with an analytic p-profile can be described by equation (2.4) with a profile $h(w)$ defined by*

$$\ln h(w) + \sum_{k=1}^{\infty} \frac{1}{k} \lambda_k \left[h^k(w) - 1 \right] = w \tag{2.10}$$

with coefficients λ_k subject to the normalization condition (2.9).

Proof. Insert the series expansion (2.8) in the formal solution of the equation of growth (eq. 2.7) and integrate term by term. □

Functions defined by equation (2.10) will be called *hindering functions*. While they cannot be determined analytically from this implicit definition, given a set of coefficients $\{\lambda_k\}$ the corresponding $h(w)$ can be computed at any w with a suitable numerical procedure; the Newton method proved both effective and efficient.

The set $\{\lambda_k\}$ and the parameters γ, Q_τ and τ uniquely pin down every hindered growth processes with analytic p-profile. To understand how each parameter affects the growth pattern, it is useful to examine the behavior of $\hat{Q}(t) \equiv \ln Q(t)$ as a function of t:

$$\hat{Q}(t) = \hat{Q}_\tau + \ln h\Big(\gamma(t - \tau)\Big)$$

The pair $(\tau, \hat{Q}_\tau)$ centers the origin at a well defined feature of the growth pattern. By scaling the time axis, γ controls the slope of $\hat{Q}$: the time derivative is $d\hat{Q}/dt = \gamma\, h'/h$; for the same profile $h(w)$, i.e., a given set $\{\lambda_k\}$, the higher is γ the steeper is $\hat{Q}$. Finally, the set $\{\lambda_k\}$ determines the concavity of $\hat{Q}$. Panel (a) in Figure 1 shows single term hindering functions, those where the set $\{\lambda_k\}$ is zero for all but a single k. The higher the order k, the more concave the function $\hat{Q}$.

Hindered growth processes described by analytic p-profiles are subject to a strict upper bound.

Lemma 1. *$h' < 1$ whenever p is an analytic function.*

Proof. Equation 2.6 gives $h' = h/(1+p)$. When $w < 0$, $h < 1$; and because $p \geq 0$, $h' \leq h < 1$. When $w > 0$, $q > 1$ so $q^k > q$ for all $k > 1$ and $p \geq \sum_{k=1}^{\infty} \lambda_k q = q$. Because $h(w) = q(t)$, $p \geq h$ and $h' \leq h/(1 + h) < 1$. □

We refer to the case $\lambda_1 = 1$, where $p(q) = q$, as *linear hindering*. It is the only case to produce $h' \to 1$ in the limit $t \to \infty$.

Theorem 2. *A variable Q corresponding to a hindered growth process with an analytic p-profile cannot increase by more than γQ_τ per unit time.*

Proof. Follows directly from $h' < 1$ and the equation of growth

$$\frac{dQ}{dt} = \gamma Q_\tau h'(w). \tag{2.11}$$

The upper bound γQ_τ is approached asymptotically for linear hindering only. □

2.2. *Single-Term Hindering*

The simplest hindering functions are obtained when all but one λ_k vanish; that is, $\lambda_k = 1$ for some single k and 0 otherwise. Denote by h_k the corresponding single-term hindering function (sth) of order k, then $p = h_\mathrm{k}^k$ and

$$\ln h_\mathrm{k}(w) + \frac{1}{k} h_\mathrm{k}^k(w) = w + \frac{1}{k} \qquad k = 1, 2, \ldots \tag{2.12}$$

Figure 1 shows in its top panel a sample of sth functions together with two limit cases. Equation (2.12) implies that $h_\mathrm{k}(w) = [h_1(kw)]^{1/k}$; every h_k can be determined from h_1, the linear-hindering function. With this relation it is straightforward to show that

$$\lim_{k\to\infty} h_\mathrm{k}(w) = \begin{cases} e^w & w \le 0 \\ 1 & w > 0 \end{cases} \tag{2.13}$$

In addition, while the sth ensemble is defined only for $k \ge 1$ it can be extended by taking the $k \to 0$ limit. In this limit $(z^k - 1)/k \to \ln z$ for all $z > 0$ (from L'Hôpital's rule), therefore $h_\mathrm{k}(w) \to e^{w/2}$. The exponential function can thus be considered the h_0 member of the ensemble of hindering functions, the only one with a constant growth rate.[3]

As is evident from the h_k plots, $\ln h_\mathrm{k}$ becomes more concave with increasing k; at $w = 0$, its Arrow-Pratt concavity measure is $-(\ln h_\mathrm{k})''/(\ln h_\mathrm{k})' = k/4$. It is also evident that the point $w = 0$, where $h = 1$ for all hindering functions, marks a transition between two distinctly different time variations. In the domain $w < 0$, the w-dependence of the left-hand side of eq. 2.12 is dominated by the logarithmic term, yielding $h_\mathrm{k}(w) \simeq \exp(w + \frac{1}{k})$. We refer to this as *unhindered growth*. In contrast, the logarithmic term can be neglected when $w \gg 1$, and $h_\mathrm{k} \gg 1$, yielding $h_\mathrm{k} \propto w^{1/k}$, namely $Q \propto t^{1/k}$ — the larger is k, the slower the growth. Instead of the universal exponential variation of unhindered growth, the asymptotic behavior of *hindered growth* is k-dependent, as is also evident from

[3] Elitzur (2022) stated erroneously that $h_0(w) = e^w$ instead of $e^{w/2}$

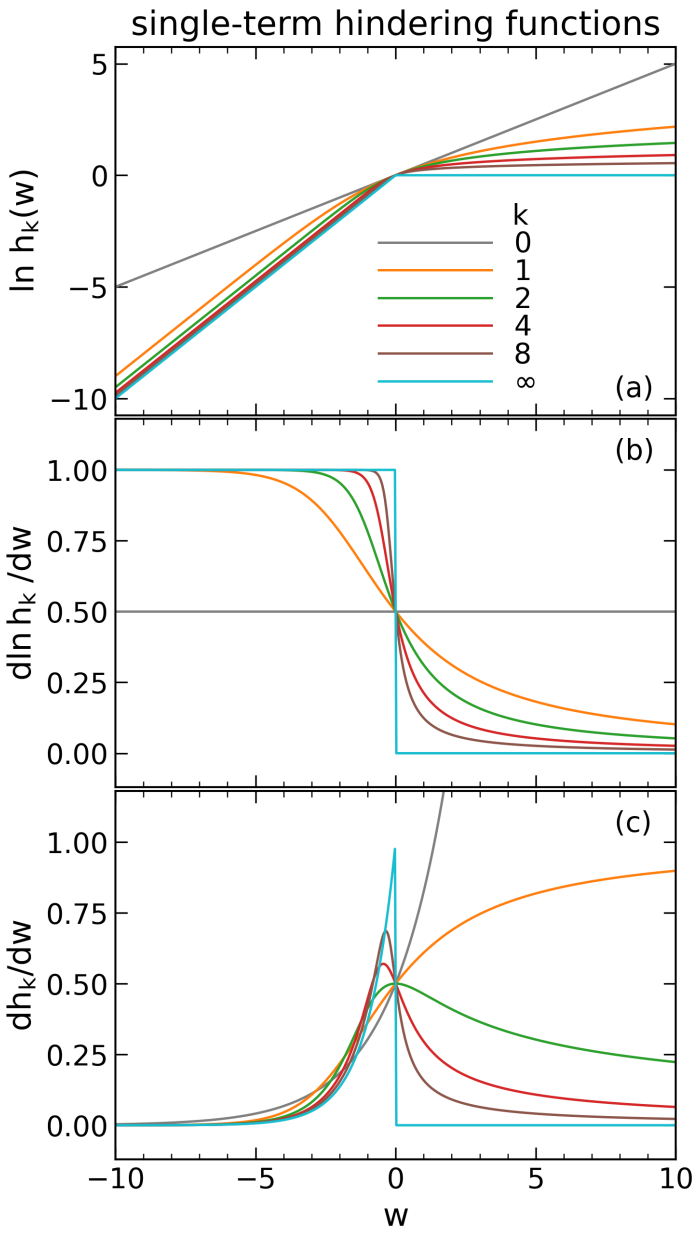


Figure 1: *Top*: A sample of single-term hindering functions h_k (eq. 2.12; see text for the limits $k \to 0$ and ∞). *Mid*: The logarithmic derivative ($= g/\gamma$). *Bottom*: The derivative $h_k{}'$.

the figure. The transition from unhindered to hindered growth becomes more abrupt with increasing k.

The mid and bottom panels of Figure 1 show the associated growth rates and derivatives. The bound derived in Lema 1 is evident in the plots of h_k' (panel c). For all $k > 1$, $h_k{}'$ peaks at $w = -\frac{1}{k}\left[\ln(k-1) + \frac{k-2}{k-1}\right]$ at a value of $k^{-1/k}\left(1 - \frac{1}{k}\right)^{1-1/k}$. The peak of h_2' is $\frac{1}{2}$ at $w = 0$. As k increases, the peak location first moves to the left, then back toward $w = 0$; the leftmost peak, when $k = 4$, is at $w = -0.441$. As $k \to \infty$, the peak of h_k' is slowly approaching 1 at $w = 0$. Denote that peak value μ_k. From equation 2.11, $\dot{Q}$ will reach a peak of $\dot{Q}_{\text{peak}} = \mu_k \gamma Q_\tau$. The ratio $\dot{Q}(t)/\dot{Q}_{\text{peak}}$ of the actual production rate of Q to its peak value is an efficiency factor known as the *capacity utilization rate* — when it reaches 1, the system grows utilizing its resources to the fullest.

2.3. Multi-Term Hindering

When the series expansion of p (eq. 2.8) contains more than one term, the growth process behaves very differently for finite and infinite number of terms:

Theorem 3. *When the set $\{\lambda_k\}$ is finite, $Q \to \infty$ as $t \to \infty$. When the set is infinite, growth is bounded by the carrying capacity $K = \kappa Q_\tau$ with $\kappa > 1$.*

Proof. When the set $\{\lambda_k\}$ contains a finite number of coefficients, p is finite for every finite q so $g > 0$ (eq. 2.5) and growth continues indefinitely; the growth rate vanishes only in the limit $q \to \infty$ — growth is unbounded. On the other hand, since all non-vanishing λ_k are > 0, when the set is infinite there is some finite $\kappa > 1$ such that $p(q)$ diverges for all $q \geq \kappa$. Thus the growth rate vanishes when $q \to \kappa$ and growth is bounded by the *carrying capacity* $K = \kappa Q_\tau$. □

Corollary 3.1. *The logistic is a hindering function, described by eq. (2.10) with $\lambda_k = 2^{-k}$ for $k = 1, 2, \ldots$. Its carrying capacity is $K = 2Q_\tau$.*

Proof. See Appendix A.2. □

Some additional discussion of the properties of hindered growth processes and the hindering functions is presented in Appendix A.3 and Appendix A.4.

3. US GDP 1790–2023

The *Measuring Worth* project compiles historical US population and GDP annual data extending back to the 18th century, and we use its latest release (Williamson, 2024). We first verify that US GDP data meet the first two conditions of Definition 1. The third condition is unverifiable with finite data. We

check for a monotonically increasing trend with the Mann-Kendall test (hereafter MK test; Kocsis et al. 2017). This non-parametric test determines whether or not a time series has a monotonic trend without assuming that its data points are distributed according to any specific rule; in particular, there is no requirement that they be normally distributed. The null hypothesis is no trend in the time series, in which case the test statistic Z is distributed according to $N(0, 1)$; that is, a dataset with MK statistic Z is Z standard deviations away from the null distribution center with a positive (negative) Z indicating a monotonically increasing (decreasing) trend. The US GDP dataset has $Z = 22.6$ (the maximum possible for 234 points is 22.8), overwhelmingly establishing the presence of a monotonically increasing trend.

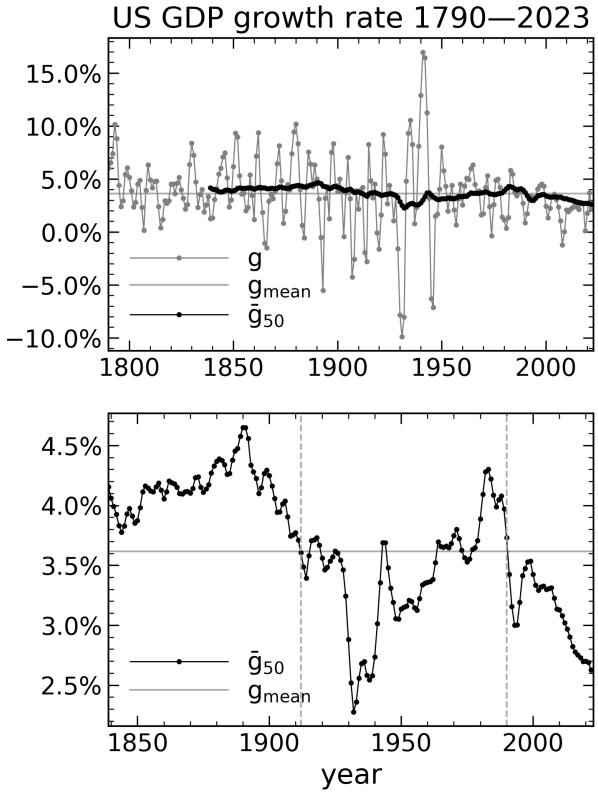


Figure 2: *Top*: The GDP annual growth rate g, its mean g_{mean} and 50-year rolling average $\bar{g}_{50}$. *Bottom*: The variation of $\bar{g}_{50}$ shows 3 distinct epochs, delineated with vertical dashed lines: until 1911, $\bar{g}_{50} > g_{\text{mean}}$ every year; from 1991 onward, $\bar{g}_{50} < g_{\text{mean}}$; in between it fluctuates both above and (mostly) below g_{mean}.

The other requirement for hindered growth is that the growth rate be either

constant or monotonically decreasing. We calculate the time series of annual growth rate by numerical differentiation. Linear regression of the growth rate as a function of time yields a downward slope, implying decline of the annual growth rate at the rate of $8.63{\cdot}10^{-3}$ percent per year. The p-value of this result is $4.21{\cdot}10^{-3}$. The MK statistic of the growth-rate time series is $Z = -3.60$, rejecting the null hypothesis of no downward monotonic trend with a p-value of $1.61{\cdot}10^{-4}$. Both linear regression and the MK test independently establish the presence of a declining trend in the growth rate with a high degree of confidence.

The plot of g in the top panel of Figure 2 shows a wildly fluctuating quantity. The time-series mean is $g_{\text{mean}} = 3.62\%$ and its standard deviation is 3.39%, almost as large as the mean (the minimum is −9.9%, reached in 1931, the maximum 17% in 1941). The growth rate declining trend, decisively established by both linear regression and the MK-test, is hard to discern because of the large fluctuations. Some direct evidence for this decline can be teased out by averaging over the fluctuations. The mean growth rate is 4.15% for the first 50 years of the time series, but only 3.47% for the rest of it; the mean for the final 50 years is just 2.63%. In addition to g, the figure also shows its rolling average $\bar{g}_{50}$, the mean of the 50 previous years, beginning with 1839; this time series has a greatly reduced standard deviation of 0.53%. An expanded y-scale in the bottom panel, where g is omitted, provides a clearer view of the variation of $\bar{g}_{50}$, illustrating the downward trend of the growth rate: In its first 73 years, from 1839–1911, $\bar{g}_{50}$ never dipped below g_{mean}. During the following 79 years it fluctuated both above and below g_{mean}, rising above it 31 times (39%). After 1990 $\bar{g}_{50}$ never rose to the level of g_{mean}.

3.1. Hindering Detrending

We decompose the output time series, Y_t, into its two components—growth and business cycle—as follows:

$$Y_t = Q_t + X_t Y_t \tag{3.1}$$

Here growth, Q_t, is given by a hindering function, and the dimensionless fluctuations X_t follow an ARMA process

$$\phi(L)X_t = \theta(L)\varepsilon_t \tag{3.2}$$

where $\phi(L)$ and $\theta(L)$ are lag polynomials of arbitrary length and $\varepsilon_t \sim$ is white noise.[4] To estimate the parameters that govern the evolution of output we have

[4]It is common practice to fit the log of Y_t rather than Y_t. In Appendix B.1 we discuss this approach and point out its shortcomings.

Table 1: PARAMETER ESTIMATES

	(a) data	(b) training set
γ (def. 1.3; percent year^{-1})	3.95 (0.02)	3.95 (0.03)
Q_τ (eq. 2.2; \$ trn)	26.61 (2.84)	27.34 (5.24)
τ (eq. 2.2; years)	240.52 (2.49)	241.42 (4.57)
calendar year of τ	2030	2031
AR(2), lag 1	1.33 (0.04)	1.33 (0.04)
AR(2), lag 2	−0.37 (0.04)	−0.37 (0.04)

Notes: Parameters for the full dataset, 1790–2023 (column a), and training set, 1790-2003 (column b; see §4.1). In both cases the trend component is $k = 1$ single-term hindering, the ARMA component is `AR(2)`. Numbers in parentheses are the standard errors.

to find the hindering function that describes Q_t and the ARMA components of X_t.

The detrending process involves three steps, whose details are described in Appendix B.2. First, we search the ensemble of hindering functions for the best fitting trend function that minimizes $\sum X_t^2$. Following the discussion in §2, the trend is described analytically by the combined equations 2.4 and 2.10. Next, the resulting fluctuations are ARMA-fitted in a systematic search for the order $(\mathtt{p}, \mathtt{q})$ that minimizes the Bayesian information criterion (BIC). Finally, the hindering function determined in the first step and the order $(\mathtt{p}, \mathtt{q})$ in the second are combined to determine the parameters γ, Q_τ, τ and ARMA lag coefficients that minimize $\sum \varepsilon_t^2$.

Table 1 lists the parameter estimates for the US GDP from 1790–2023. The mean magnitude of the deviation of model from data is 2.75%. The trend component is $k = 1$ single-term hindering (eq. 2.12). It is selected by the F-test over the exponential function with overwhelming probability — the null hypothesis of no hindering term is rejected with a p-value of $1.11 \cdot 10^{-16}$. The fluctuation component is `AR(2)`; notably, the best fit does not have an `MA` component. The GDP has not yet reached the hindering point (eq. 2.2) and is expected to do that only in 2030 — the data sample only negative values of the independent variable w (eq. 2.3) in the range $[-9.51, -.30]$. Nevertheless, hindering is already impacting significantly the growth of US GDP, with the underlying annual growth rate down to 2.11% from the initial $\gamma = 3.95\%$.

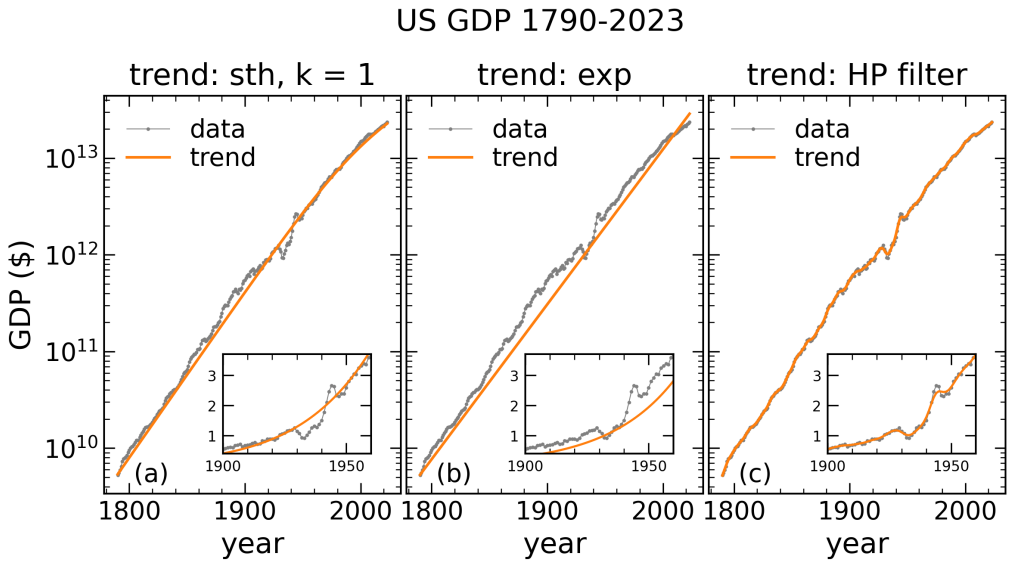


Figure 3: US GDP ($ = 2020 USD) from 1790–2023; the inset in each panel zooms in on the 1900–1960 period with a linear y-axis in $ trillions. In addition to the data, each panel shows the trend component of a different model, as labeled; the arma component is omitted to reduce clutter. Trends based on analytic hindering functions are shown for (*a*) $k = 1$ single-term hindering (eq. 2.12) and (*b*) the exponential. Panel (*c*) shows the trend component of the Hodrick and Prescott (1981) decomposition method (§3.2).

For further insight into the significance of hindering, the left and mid panels of Figure 3 (panels *a* and *b*) show the data together with the trend components of the models based on the linear-hindering and exponential functions, even though the latter is decisively rejected by the F-test. Thanks to the logarithmic scale, a downward concavity is clearly discernible in the data. This concavity is captured by the hindering fit, but not by the exponential one, which can only select the best fitting straight line through the data (the selection is affected by the coupled ARMA component, which is not shown for clarity). In each panel the inset zooms in on the first 60 years of the 20th century, a period that includes the Great Depression and the World War II production boom. In spite of these major shocks to the US economy, its underlying growth trajectory reverts to the same trend it had before the upheavals of the 1915–1950 turbulent years, closely following the linear hindering fit.

Figure 4 shows the GDP growth rate, providing additional perspective on the differences between the linear hindering and exponential trends. Prior to 1950, the growth rates from the two models are almost indistinguishable. However,

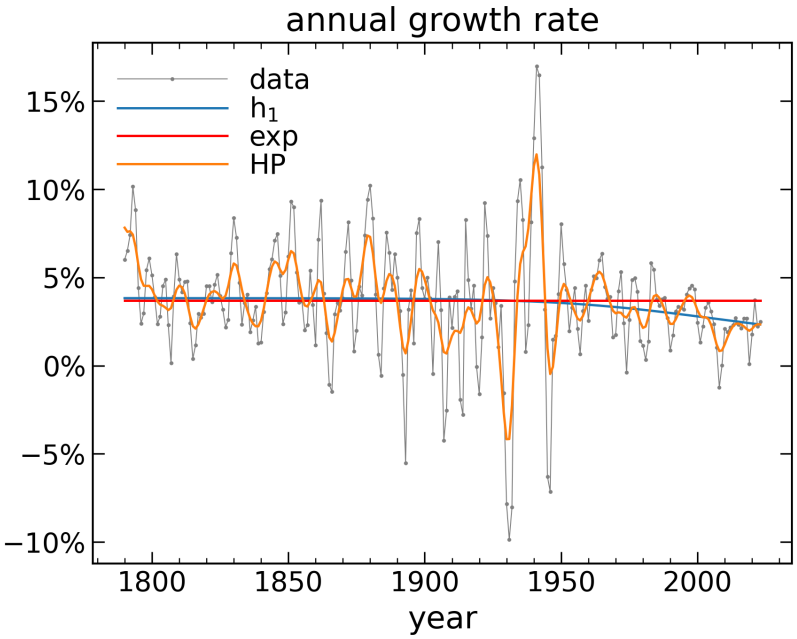


Figure 4: Annual growth rate of US GDP from 1790–2023. In addition to the data, the growth rates of the various trends displayed in Figure 3 are also shown: linear hindering (labeled h_1), exponential (exp), and the Hodrick and Prescott (1981) decomposition (HP).

the subsequent declining trend, evident in the data (see also Figure 2), is properly captured by the linear-hindering trend while the exponential maintains its constant value. Our model implies that the underlying annual growth rate of US GDP follows the simple numerical expression

$$g = \frac{3.95\%}{1 + \text{GDP}/\$26.6\text{ trn}} \tag{3.3}$$

Figure 5 shows the fluctuations from each detrending method. Plotted is the fluctuation z-score; the mean and standard deviation (std) of each distribution are listed in Table 2. Detrending aims to yield stationary fluctuations, and stationarity can be checked with the augmented Dickey–Fuller (ADF) test. The test null hypothesis, that a unit root is present in the time series, is rejected decisively for the hindering fluctuations with a p-value of 0.006. In contrast, the p-value of the exponential fluctuations is 0.252, thus the presence of a unit root cannot be rejected. Nelson and Plosser (1982) proposed the unit root hypothesis. Based on exponential detrending of the US GDP, they reported being unable to reject the presence of a unit root and suggested that GDP growth trend itself is stochastic,

described by non-stationary random walk with drift. Our findings agree with theirs — we, too, are unable to reject the unit root hypothesis with exponential detrending. However, the hypothesis is rejected decisively when detrending properly employs the linear hindering function instead of the exponential.

Table 2: FLUCTUATION MOMENTS

	h_1	exp	HP
mean	0.082 (0.0087)	0.23 (0.0093)	−0.00038 (0.0018)
std	0.13 (0.0062)	0.14 (0.0066)	0.028 (0.0013)

Notes: Mean and standard deviation of the distributions of X_t (eq. 3.1) for the hindering, exponential and HP decomposition analyses. Numbers in parentheses are the sample standard errors (ste).

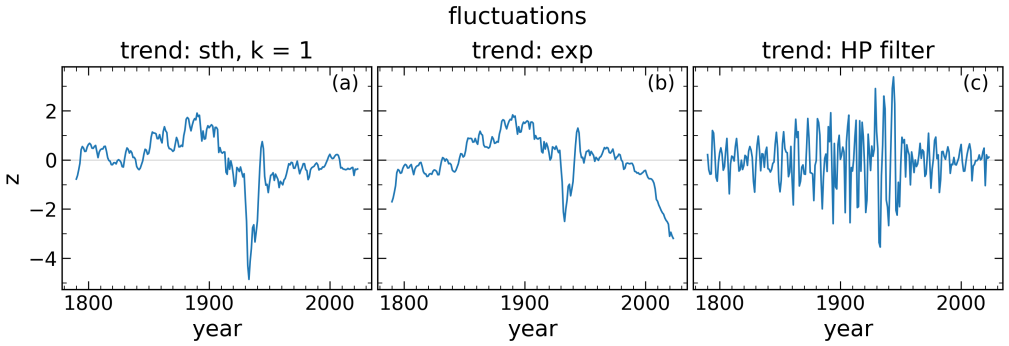


Figure 5: The fluctuation component of the three decomposition methods, labeled as in Figure 3. The plotted quantity is the fluctuation z-score.

In addition to failing the ADF stationarity test, the fluctuations from exponential detrending produce results at odds with the historical record. The mid panel (*b*) of Figure 5 shows that more than 70 years ago, GDP growth has entered a downward fluctuation that is yet to show signs of reversing itself. This fluctuation is necessary to compensate for the departure of the growth rate declining trend from the constant value it has in exponential decomposition (Figures 2, 4). In its length and severity, this fluctuation greatly surpasses the Great Depression, contrary to economic data; for example, in no time during the allegedly ongoing fluctuation did the unemployment rate approach levels reached during the Great Depression. This conflict with the historical record helps explain why the Nelson

and Plosser (1982) unit root hypothesis is no longer viable. While in 1982 it was possible to suggest that the underlying growth rate is constant and that its downward trend is only a temporary fluctuation soon to be reversed in a regression to the mean, more than 40 years later this is no longer plausible.

In contrast with the unrealistic consequences of exponential detrending, the fluctuations from the hindered growth decomposition are consistent with economic history: The Great Depression stands out as the most extreme fluctuation of any kind; the Great Recession and the oil shocks of the 1970s produced downward fluctuations of comparable magnitude; the boom of the late 1990s and early 2000s induced a similar fluctuation in the opposite direction.

3.1.1. Robustness

Hindering detrending is premised on growth processes having two structured components: one deterministic, the trend, one stochastic, the fluctuations. By the very nature of stochastic processes, a given set of parameters allows infinitely many realizations, and those can be combined with a single trend to generate an infinite ensemble of growth processes, all with identical model parameters. Suppose we combined purely exponential growth with different ARMA processes generated from the same parameters. Could a result of such calculation display a hindering trend? If that were the case it would be impossible to separate exponential from hindered growth because we would have constructed a hindered growth process out of a purely exponential trend.

To asses the probability of such spurious hindering we constructed a sample of a million mock datasets in Monte Carlo simulations (Appendix B.3). The probability to display hindering by a growth process generated from fluctuations around an exponential trend with the parameters of our best fit is $7{\cdot}10^{-5}$. The chance occurrence of the MK statistic of the US GDP growth rate is one in 52 million. The possibility that the US GDP growth rate is constant and only appears to be decreasing because of chance fluctuations can be dismissed as highly unlikely.

3.2. Hodrick-Prescott Decomposition

The Hodrick and Prescott (1981, 1997) decomposition has become a standard technique for separating a time series into its trend and cyclical (fluctuation) components. While drawbacks to this method were pointed out, it remains widely adopted in academic research, policy studies, and analysis by private sector economists (Harvey and Jaeger, 1993; Hamilton, 2018). The method employs the operator $H = (I - L)^2(I - L^{-1})^2$, where I and L are, respectively, the unit and lag operators (Baxter and King, 1999). The H-operator performs a 5-year rolling

average with weight coefficients extending symmetrically two years into the past and the future. The trend component of the time series is obtained by operating on it with $T = (I + \lambda H)^{-1}$, where λ is a smoothing parameter, the cyclical component with the operator $I - T$.

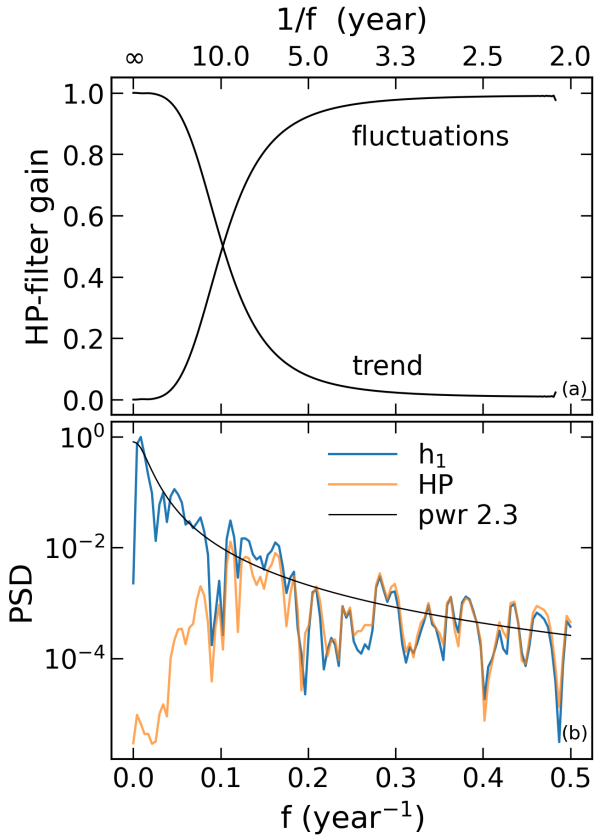


Figure 6: Spectral properties in terms of Fourier frequency f of the HP and hindering decompositions. *Top*: The HP-filter gains of the trend and fluctuation (cyclical) components, as labeled. *Bottom*: Power spectral density (scaled units) for the hindering (h_1) and HP fluctuations, as labeled. A power-law fit to the h_1 PSD, declining as $f^{-2.3}$, is also shown (see text for details).

We apply the HP decomposition with $\lambda = 6.25$, recommended by Ravn and Uhlig (2002) for annual data. Since this method is usually applied to log of the data (e.g. Razzak, 1997) we calculated the trend both directly from the data and by exponentiating the trend component from analysis of log of the data. The results from the two calculations differ by less than 1%. The fluctuations are then calculated from eq. 3.1, their best fit is `ARMA(2, 2)`. The derived HP trend is shown in the right panel (c) of Figure 3, the growth rate in Figure 4 and the fluctuations

in the right panel (c) of Figure 5. Each plot provides its own perspective on the differences between the HP and hindering decomposition methods. The hindering trend, shown in the left panel (a) of Figure 3, delineates the GDP overall trajectory — exponential growth with slowly, monotonically decreasing growth rate (Fig. 4; eq. 3.3). In contrast, the HP trend and its growth rate closely trace the data, smoothing out structure with time scale up to ∼3 years but leaving all longer-term variability largely intact. The insets in Figure 3 amplify this point. In stark contrast with linear hindering, even the Great Depression and the World War II production boom are parts of the HP growth trend instead of standing out as the large deviations from norm that they are. These points are further demonstrated by the fluctuation plots in Figure 5. The HP filter fluctuations are both shorter and more uniform than those of hindering and exponential detrending.

The results of HP decomposition reflect the frequency variation of the Fourier representation of the HP filter gains (ratio of emergent to input signal), shown in the top panel of Figure 6. The trend and fluctuation gains cross at frequency $f = 0.1$ year^{-1} (cycle period 10 years), thus an input signal at that frequency will be shared equally between the two components. Higher frequency signals are assigned preferentially to the cyclical component, lower ones to the trend. As a result, the power spectral density (PSD), the distribution of power among Fourier components, of the HP fluctuations is nearly identical to that of the hindering decomposition at frequencies higher than ∼0.1 year^{-1} but virtually devoid of all lower frequencies (Figure 6, bottom panel). Neither PSD shows significant structures. The apparent peak around f ∼0.1–0.15 in the HP fluctuations PSD, the signature of periodic signals, is an artifact of the HP filtering. The hindering PSD is free of such artifacts and can be considered a more faithful representation of the intrinsic fluctuation spectrum. While highly fluctuating, it gives no indication of outstanding peaks and can be fitted with a power-law variation of the form $a/[1 + (f/f_0)^p]$. The bottom panel shows the fit with $p = 2.3$.

The HP filter smoothes the data sharp peaks and troughs, assigning the "trend" label to time variations that lasts longer than ∼10 years. As noted by Razzak (1997) and in contrast with the hindering analysis, the HP decomposition is a smoothing, not a detrending technique.

4. Forecasting

Our hindering model can be used for forecasting beyond the end of the time series. Figure 7 shows the model forecast for the 20 years past the end of the data. In addition to data and model, the figure also shows the model 95% confidence

interval CI(95). It is determined, together with the model standard error, by sampling the parameter space with multivariate normal distribution centered on the model optimal parameters and using its covariance matrix. Such sampling is justified by the number of data points (234) greatly exceeding the number of model parameters (5), implying asymptotic approach to Gaussian regime by the posterior probability distribution function for fitted model parameters (Press et al., 2007; Ivezić et al., 2014; Brockwell and Davis, 2016). As a check on the sample size we calculate for every point of the time series the absolute deviation of the sample mean from the model derived from the optimal parameters. The sampling is acceptable when the mean of all deviations is smaller than a prescribed accuracy; if the accuracy is not met, the sample size is increased. Sampling 1,000 times was found sufficient for better than 2% accuracy.

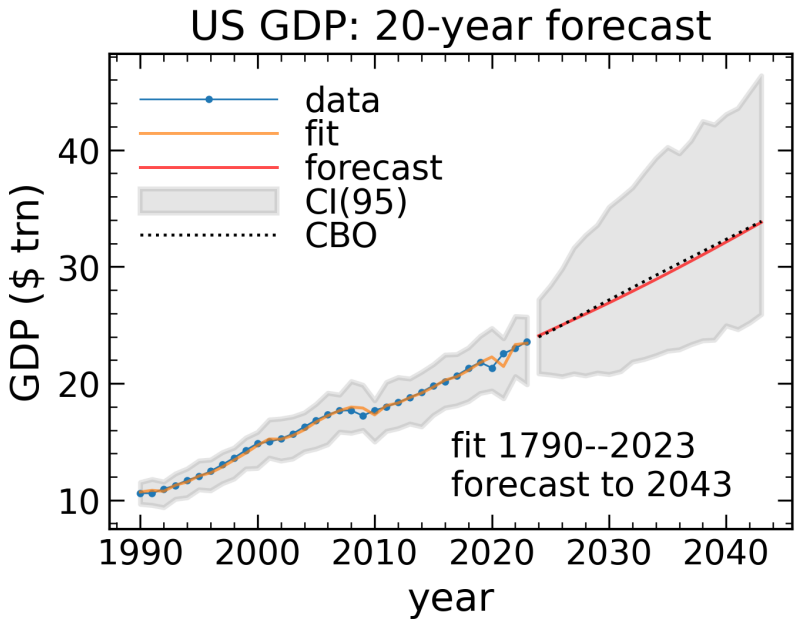


Figure 7: Twenty-year forecasting for US GDP. The data are shown together with the hindering model, its 2024–2043 forecast and the model 95% confidence interval. Also shown is the most recent US Congressional Budget Office projection, labeled CBO. Results, shown starting in 1990, are in trillions of 2020 USD.

Also shown is the long-term projection published by the US Congressional Budget Office (CBO)[5] in March, 2024. Our forecast and the CBO projection are

[5]CBO report #59711, *The Long-Term Budget Outlook: 2024 to 2054*

within 1.2% of each other. For 2043 our model is forecasting a US GDP of $33.8 trillion, the CBO projection is $33.9 trillion. The underlying annual growth rate is forecasted to decline from 2.11% in 2024 to 1.75% in 2043; the CBO projection for 2043 is 1.56%. The CBO projections are obtained from elaborate modeling, with a multitude of assumptions and a large number of parameters, entirely unrelated to ours. The close agreement between the results from two such different approaches is encouraging.

4.1. Training-Set Test

To assess the reasonableness of forecasting s time steps into the future we adopt the training-set approach from machine learning (Ron, 1998): Remove the last s members of the time series; model the truncated dataset; use the resulting model to compute s time steps past its end; compare this prediction[6] with the removed data.

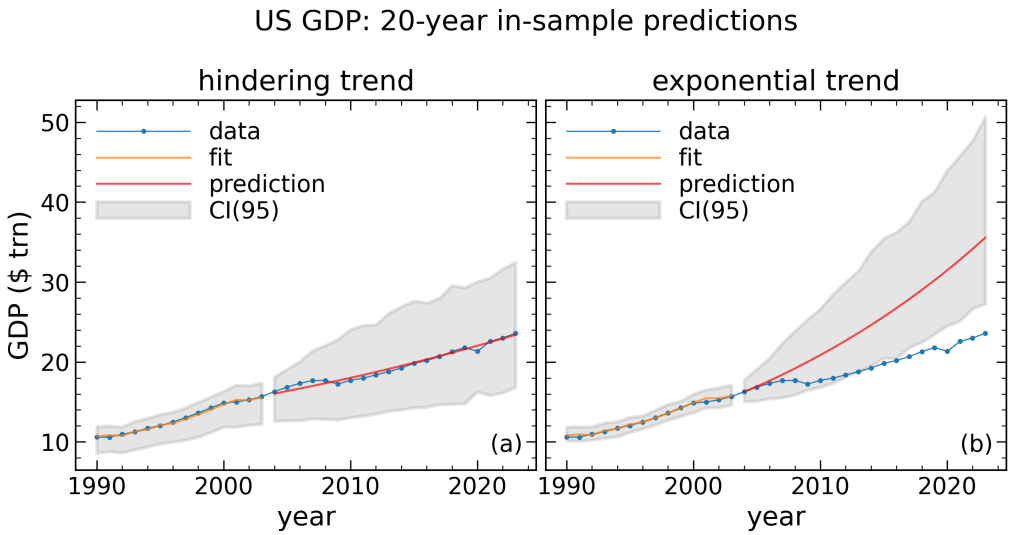


Figure 8: Training-set results for US GDP. The data are shown together with models for the 1790–2003 truncated dataset and their predictions for 2004–2023; the model 95% confidence intervals are shown in shaded gray. Panel *(a)* shows the model based on the best-fitting hindering trend, panel *(b)* does the same for exponential trend. Results, shown starting in 1990, are in trillions of 2020 USD.

Figure 8 shows our US GDP 20-year training-set analysis. The truncated data from 1790–2003 are fitted, and the resulting model predictions for 2004–2023 are compared with the removed data points. The best-fitting model of the

[6]We adopt the terminology of in-sample prediction and out-of-sample forecast.

training set, shown in the left panel (*a*), is comprised of a linear hindering trend and AR(2) fluctuations, just like the full dataset (§3). Moreover, Table 1 shows that the parameters of the two models are nearly identical, both trend and ARMA components. The main difference is the larger uncertainty in the location of the hindering point for the training set. The reason is that the hindering impact is appreciable only during the final ∼70 years of data (fig. 4) and of those, the training set removes the last 20, which are the most impactful.

The model successfully predicts every data point from 2004–2023. The mean absolute deviation of prediction from data is 1.7%. For comparison, the right panel (*b*) shows the training-set prediction of the best-fitting model based on exponential trend. This model describes successfully the truncated dataset thanks to its ARMA component, which compensates for the inadequate trend. But since out-of-sample ARMA results revert quickly to the series mean (Brockwell and Davis, 2016; Hamilton, 2020), the training-set predictions are dominated by the trend component. And because the actual growth rate is declining, the exponential model overshoots the data past the truncation point by a continuously increasing margin. Beginning in 2013, every data point is outside the lower bound of the model 95% confidence interval.

An important implication of figure 8 is that the Great Recession merely *coincided* but did not cause any structural changes to the underlying growth of the US economy. We estimate the same hindering function using data that predates the Great Recession, and we show, in fact, that the slowdown in growth was already under way when the Great Recession hit. Since long-term growth is determined by supply, rather than demand, factors, our findings are not consistent with the secular stagnation view that the unique demand factors that arose from the Great Recession caused the slowdown in the growth trend.

4.2. An Upper Bound on US Economic Growth

The US GDP rose from $18.8 trillion in 2013 to $23.6 trillion in 2023, an average annual increase of $0.48 trillion (the model result is $0.47 trillion). From Theorem 2 (see §2), hindered growth with analytic p-profile has an upper bound — the growing quantity can increase asymptotically at most linearly with a slope of γQ_τ. Our modeling shows that the growth trajectory of the US economy is linear hindering with γQ_τ = $1.05 trillion per year (Table 1). Therefore *adding every year $1.05 trillion to the GDP is the upper bound on the growth of the US economy.* This asymptotic limit is still far from being reached. The current value of the intrinsic independent variable w (eq. 2.3) is −0.31, which Figure 1 shows is at early stages of deviation from exponential growth.

5. Fluctuation Analysis

We now focus on the fluctuations around the trend. In remaining agnostic about the trend, our methodology also remains agnostic about the dynamic structure of the fluctuations. In contrast to filtering methods, which constrain the frequencies that fluctuations can contain, our approach recovers fluctuations as the difference between the data and the trend. Rather than impose structure on the data, we let the data reveal that structure to us.

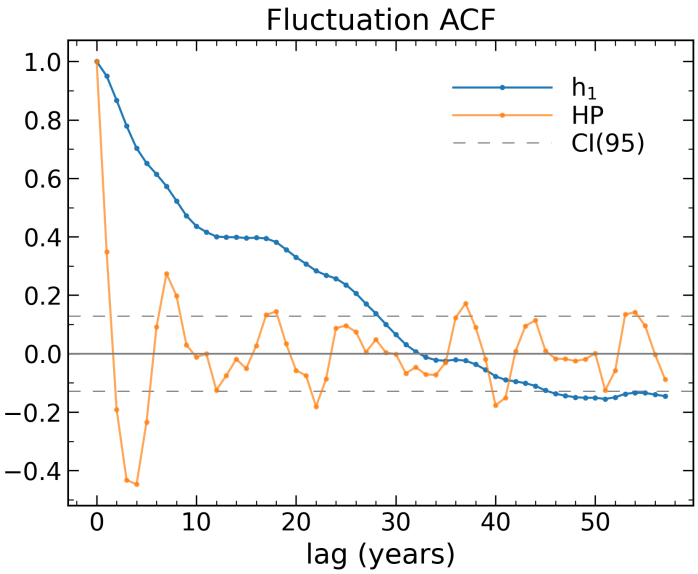


Figure 9: First 58 lags of the autocorrelation functions (ACF) of the fluctuations around the hindering and HP trends, as labeled. Dashed lines are the 95% confidence boundaries of random noise ACF.

Consider the fluctuations around the trend components of the hindering and HP decompositions. When fluctuations are compatible with random noise at the 95% confidence level, no more than ~3 of the first 58 autocorrelations beyond zero lag fall outside the CI(95) noise boundaries (see Appendix B.4). Instead, Figure 9 shows that the fluctuations from each decomposition method contain significant, though very different, temporal structure. The ACF of the HP fluctuations drops steeply after lag zero as expected from random noise, but 15 of its subsequent values (26% of total) are outside the CI(95) boundaries. These outliers display prominent features that mirror technical details of the HP method rather than intrinsic properties. A structure between lags 2–5 protrudes below the noise lower boundary by as much as factor of 3.5. This negative correlation

arises from the HP minimization recipe, which replaces each data point with a 5-years average that includes 2 leading and 2 trailing years (§3.2). Positive correlations at 7–8 years rise to more than a factor of 2 above the noise upper boundary. They reflect a deliberate attempt to align the HP filter with presumed typical durations of business cycles through a judicious choice of the smoothing parameter λ (Baxter and King, 1999; Ravn and Uhlig, 2002). These two structures generate higher harmonics that remain mostly within the noise boundaries. Other than traces of details of the technique itself, the HP decomposition smoothes out the fluctuations around the trend component to below the noise level.

In contrast with the HP decomposition, hindering detrending does not take into considerations any temporal structure that may or may not exist in the fluctuations around the trend. Its only prior assumptions, spelled out in Definition 1 (§2), involve general properties of the growth trend. Therefore, autocorrelation analysis and ARMA modeling of the hindering fluctuations provide direct information about the stochastic variability of US GDP, uncovering time scales intrinsic to this variability. The fluctuation ACF declines slowly and crosses the noise upper boundary at lag separation of 28 years (Figure 9). That is, correlations between different fluctuations are recognizable above the noise level as long as their separation time does not exceed $t_{\text{noise}} = 28$ years. We will classify ACF $> \frac{1}{2}$ as a strong correlation because it is closer to 1, a perfect correlation, than to 0. With this classification, fluctuations are strongly correlated when separated by up to $t_{\text{strong}} = 10$ years, the lag where the ACF drops below $\frac{1}{2}$; since $t_{\text{noise}} > 2t_{\text{strong}}$, the ACF decline is steeper at smaller lag separations. Another characteristic time scale is determined from the ARMA modeling — members of an AR(p) process are directly coupled within p time steps of each other. With Δt the series time interval, the duration of direct coupling is $t_{\text{AR}} = p\Delta t$. Since AR(2) is the best ARMA fit to the fluctuations (§3.1), US GDP fluctuations are directly coupled for $t_{\text{AR}} = 2$ years.[7]

Stochastic variability can be characterized also by its impulse response function (IRF). The IRF, derivable from the ARMA fit, describes how a system would evolve after a single, brief impulse (a shock) of unit strength. By definition IRF(0) = 1, while n time steps from the shock the IRF of an AR(2) process is a linear combination of geometric terms r^n, where r is a root of the lag polynomial (Cochrane, 2005). Denote $t = n\Delta t$, then $r^n = e^{t/\tilde{\tau}}$ with $\tilde{\tau} = (\ln r)^{-1}\Delta t$ — the IRF is a superposition of exponentials and, when r is complex, damped harmonic

[7] Our analysis employs annual GDP. Quarterly data, if available, might have led to AR(6) as the best ARMA fit, in which case t_{AR} would be 18 months.

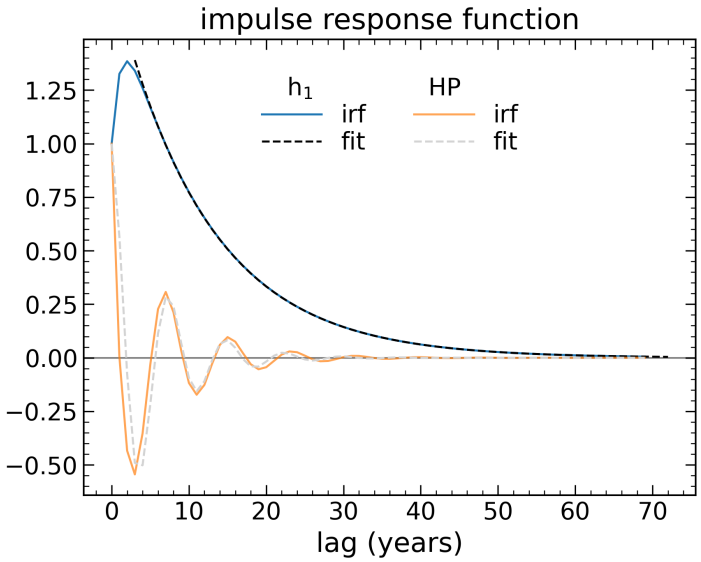


Figure 10: Impulse response functions of the ARMA components of the hindering (h_1) and HP decompositions, as labeled. Solid lines show the computed IRF, dashed lines are numerical fits (see text for details).

oscillations.

The hindering IRF, shown in Figure 10, rises to a peak of 1.38 in 2 years; notably, the rise time is $t_{\rm AR}$. The subsequent decline follows closely exponential decay, plotted in dashed line. The decay constant is $\tilde{\tau} = 11.9$ years, corresponding to one of the two roots of the lag polynomial; the decay from the other root is over in less than a year. As a result of its decline, the IRF becomes indistinguishable from zero, to within its standard error, after lag year 31. This is comparable to the estimate of $t_{\rm noise}$ when the ACF drops below the noise level. Similarly, the decay time $\tilde{\tau}$ is comparable to $t_{\rm strong}$, the ACF estimate for the duration of strong correlations.

The hindering IRF matches the generic response of a physical system in equilibrium to the impact of a sudden, brief disturbance. The rise time is controlled by the system inertia, the peak level by the ratio of the magnitude of the disturbance to that of the forces that balance each other in equilibrium. The relaxation back to equilibrium involves more detailed properties specific to the system. The HP decomposition IRF, on the other hand, shows damped oscillations (Figure 10). The plotted fit is a cosine function with a period of 7.5 years and an exponentially decaying amplitude with decay time of 6 years. As with the ACF, the HP IRF reflects technicalities of the decomposition method, not intrinsic proper-

ties of the dataset.

5.1. The Business Cycle

Comin and Gertler (2006) note that business cycles longer than suggested by Burns and Mitchell (1946) do exist but have been missed because conventional filters tend to sweep longer oscillations into the trend.[8] They introduce a class of medium-term business cycles with a high-frequency component corresponding to the original Burns & Mitchell definition and a medium-frequency one extending all the way to 50-year cycles. Although the upper end of cycle length may just reflect the time span of the datasets in their study (all are between 1948 and 2001), the framework they propose is in accord with our hindering decomposition results. We find that US GDP fluctuations are directly coupled for two years, display strong correlations for up to ∼10 years and weaker ones for up to ∼28 years, blending with the noise at longer separations. An interesting aspect of the fluctuation ACF is the plateau at lags ∼12–18 years (Figure 9), possibly reflecting the delayed impact of certain activities on the overall state of the economy. The Comin & Gertler high-frequency component matches the domain of strong correlations, and the medium-frequency one is aligned with weaker correlations above the noise level. The classification of various economic indicators (investment, output per capita, etc.) can be expected to reflect differences in the relative strengths of their Fourier components.

5.2. ARMA Residuals

The residuals ε_t of the ARMA fits to the fluctuations (eq. 3.2) are indistinguishable from white noise for both hindering and HP decompositions (see Appendix B.4). We now study the probability distributions of those residuals. Figure 11 displays tests of normality for the two residual distributions. The top row shows the probability density function (pdf), both empirical and theoretical, in terms of the standardized residuals z; the mean and std values are listed in Table 3. The empirical pdf is plotted as both a histogram and a smooth curve obtained through non-parametric kernel density estimation (KDE) with Gaussian kernels. Also shown are plots of two normal distributions with different width parameters — σ, the data standard deviation, and σ_G, which is determined from the data inter-quartile range (Ivezić et al., 2014, eq. 3.36). The parameter σ_G is less impacted by outliers than σ, providing a more robust estimate of width; the Gaussian function has $\sigma_G = \sigma$. For both decompositions, the plots

[8] This is the case with the HP decomposition (Figure 6).

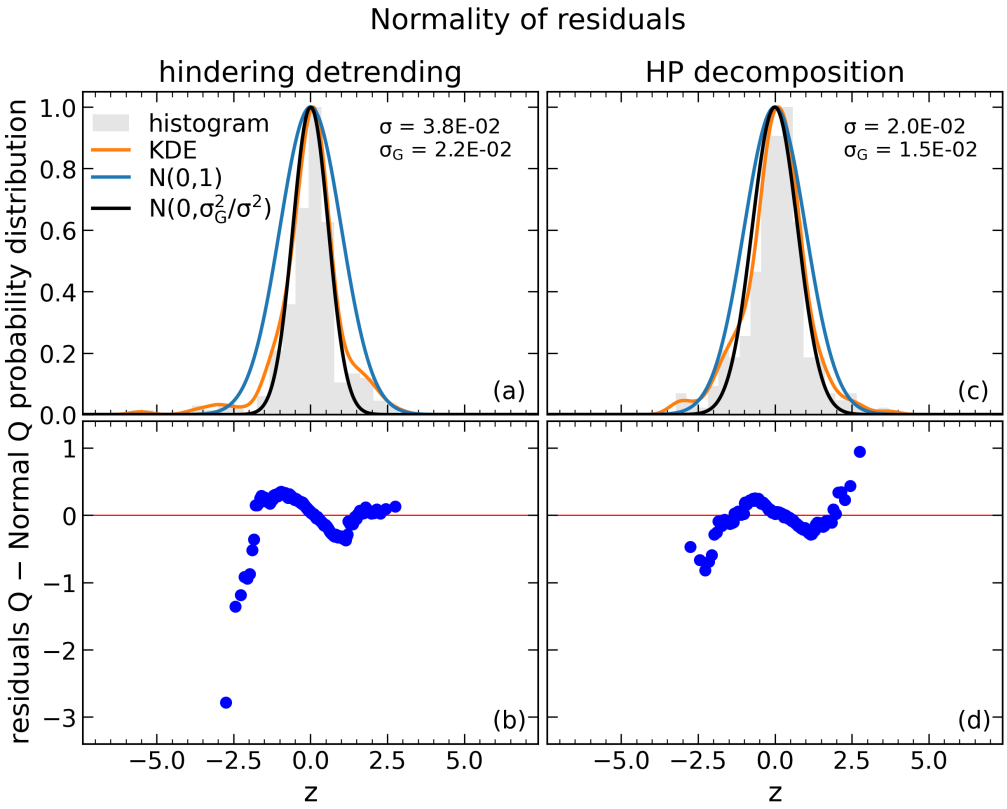


Figure 11: Normality tests for the residuals of hindering detrending (left column) and HP decomposition (right). *Top Row*: The residual pdf normalized to unity peak is shown as both a histogram and a smooth plot, labeled KDE. The variable is the residual z-score. Also shown are the normal distributions $N(0,1)$, corresponding to the residual standard deviation σ, and $N(0,\sigma_G^2/\sigma^2)$, corresponding to σ_G, derived from the distribution inter-quartile range (see text). Data drawn from a normal distribution have $\sigma_G = \sigma$. *Bottom*: Difference between the quantiles of the residual cdf and of their best-fitting Gaussians. Data drawn from a normal distribution would fall on the horizontal line at zero.

of the empirical KDE and the Gaussian with width σ_G are nearly identical up to $|z| \lesssim 2$, showing that the residuals are normally distributed in the distribution core regions. However, in both cases the two plots depart from each other in the distribution tails, $|z| \gtrsim 2$, signaling deviations from normality. These deviations manifest themselves also in the significant differences between σ_G and σ. Because of prominent non-Gaussian tails, the normal distribution with width σ is quite broader than the bulk of the pdf, especially for the hindering decomposition.

The combination of a Gaussian core and pronounced non-Gaussian tails is further illustrated in the bottom row of Figure 11, which utilizes the cumulative distribution function (cdf). Shown is the Q-Q difference which differences out the Q-Q plot, so two distributions are the same if the plot falls on the x-axis at zero. The inner regions show the hallmark of data drawn from a normal distribution — the points trace the horizontal line fairly closely. The outer regions, however, deviate significantly from normality, deviations that increase with $|z|$. The pdf and cdf both point to the same conclusion for both decomposition methods: the residuals of ARMA fits to the fluctuations are normally distributed, with significant deviations from normality at large $|z|$.

Table 3: RESIDUAL MOMENTS

	h_1	HP
mean	0.0041 (0.0025)	−0.0026 (0.0013)
std	0.038 (0.0018)	0.02 (0.00092)
skew†	−1.1 (0.16)	−0.23 (0.16)
kurtosis‡	5.4 (0.32)	1.6 (0.32)

Notes: First four moments of the ARMA residuals ε_t (eq. 3.2) of the hindering and HP decomposition analyses. Numbers in parentheses are the sample ste.
† Fisher-Pearson coefficient of skewness
‡ Fisher definition, also known as excess kurtosis (0 for normal distribution)

5.2.1. Black Swans

The skew and kurtosis listed in Table 3 show that *the hindering residual distribution is highly leptokurtic with a strong negative skew.* Notably, even though the HP decomposition is designed to trace closely the data, it too shows such deviations from normality, though not as pronounced: its kurtosis is a 5σ effect and its negative skew is a 1.4σ effect, short of statistical significance.

Table 4 lists the residuals with $|z| > 2$. Hindering detrending has 13 such outliers, the HP decomposition 12. Some HP outliers strongly conflict with the historical record. In the midst of the Great Depression, the HP decomposition implies that 1931, when g hit its all-time low of −9.9%, was a boom year with GDP growth 2.5σ *above* the underlying trend. And conversely, its GDP growth for 1941, a boom year of war production with g reaching its all-time high of 17%, is more than 2σ *below* trend. Hindering residuals have no such anomalies. The hindering outliers are well matched with historic events. In particular, the Great

Table 4: ARMA Residual Outliers

$\lvert z\rvert \geq 2$			post 1946, $\lvert z\rvert \geq 1$		
year	$z(h_1)$	z(HP)	year	$z(h_1)$	z(HP)
1861		−2.26	1948		1.70
1892		2.50	1949	−1.39	−1.52
1895	2.24	2.88	1950	1.55	
1904		−2.65	1952		−1.13
1905	2.08		1954	−1.29	−1.82
1908	−2.99	−3.22	1955	1.16	
1914	−2.84	−2.84	1958	−1.16	−1.01
1916	2.30	2.39	1959	1.14	1.06
1917	−2.42		1973		1.35
1929		3.70	1974	−1.15	
1930	−3.80		1975		−1.14
1931	−2.20	2.49	1982	−1.31	−1.04
1932	−5.54		2009	−1.04	
1934	2.54		2020	−1.22	−1.22
1938	−3.06	−3.09	2021	1.22	
1940		−3.11			
1941	2.88	−2.16			
1946	−3.46				

Years in which the z-scores of residuals obey prescribed limits. The left tabulation is for $\lvert z\rvert \geq 2$, the right one for $\lvert z\rvert \geq 1$ after 1946.

Depression years of 1930–1932 display sizeable negative consecutive outliers, reaching $z = -5.54$ in 1932. Remarkably, all 13 outliers are between 1895–1946; that is, 25% of those 52 years strongly deviate from the underlying trend, but no comparable deviations exist in either the 105 years before this period or the 77 ones after. This striking anomaly attests to the level of turbulence and upheaval in the first half of the 20th century.

Neither decomposition method produces $\lvert z\rvert \geq 2$ outliers after 1946, the largest deviations are confined to $1 < \lvert z\rvert < 2$. Thus this period might be deserving in its entirety of the name the Great Moderation. All $\lvert z\rvert > 1$ residuals during this period are listed in Table 4. Conspicuously absent from the HP results is any indication of the Great Recession, which in the hindering residuals registers as $z = -1.04$ in 2009. The hindering residuals are also consistent with the dynamics of the Covid pandemic, with two equally-sized and oppositely-signed

consecutive $|z| > 1$ in 2020 and 2021. The HP analysis does pick up the 2020 GDP drop when the pandemic struck, but not the equally fast recovery of the following year.

6. Discussion

Blanchard and Fischer (1989) identify the fundamental reason why growth decomposition is such a vexing problem — it is ill-defined. There are infinitely many ways to decompose a growth process as the sum of components characterized only as trend and stationary (cyclical). Removing this degeneracy requires a more specific definition, a challenge exacerbated by the lack of a general description of growth with diminishing marginal growth rates. This problem can be overcome by limiting exponential detrending to a period where a constant growth rate is an acceptable approximation, or segmenting the data to a number of such periods. Both approaches have been employed in various studies.

This paper presents an alternative that can be characterized as a global approach to growth decomposition. We show that placing the broad requirements of monotonicity, stability and boundedness on weakly decreasing growth rates is sufficient to define a family of log concave functions, the hindering functions (§2), which we use for the trend. Derived through rigorous mathematical solution of the generic equation of growth (eq. 2.1), the hindering functions provide a general description for growth with diminishing marginal growth rates. Such a general description would not be particularly useful if it required an unwieldy number of free parameters. However, linear hindering successfully describes 234 years of US GDP growth by modifying the exponential function with just a single free parameter (§3.1). The studied period includes waves of immigration, two world wars, the Great Depression and the transformation of the US economy from agrarian to industrial to technological. Yet with all of these massive upheavals, the GDP growth trajectory traces a single smooth function with the growth rate following a simple analytic expression (eq. 3.3). Explaining this remarkable simplicity and persistence is a challenge for macroeconomic theory.

The fluctuations around the trend are modeled with the ARMA technique. Searching for the best-fitting parameters for every possible combination of hindering function and ARMA process is unrealistic. Instead, our solution method, described in Appendix B.2, adopts a perturbation approach, one of the earliest methods of Newtonian mechanics (Bender and Orszag, 1999).[9] This approach is self-consistent if the fluctuations are a small perturbation on the trend. Indeed,

[9] Brillinger (2001) advocates a somewhat similar approach to time series analysis.

with a fluctuation std of 0.13 (Table 2), the stochastic component is only 13% of the growth process while the growth trend, at 87%, is the leading component. A similar conclusion is reached from an entirely different estimate: the fluctuation ACF drops below the noise upper boundary at lag separation of 28 years (Figure 9), at which point the fluctuation PSD is down to 10% of its peak value (Figure 6, bottom panel). The PSD of higher frequencies, with cycle periods shorter than 28 years, is smaller still, decreasing as $f^{-2.3}$, thus the fluctuations do not contain more than ~10% of total power. The agreement between these two independent estimates is an indication of the self-consistency of our analysis. The small fraction of the stochastic component creates serious difficulties for business cycle studies. The trend is such a dominant component that relatively small detrending errors can be magnified many times in the resulting fluctuations. The strong conflicts between the exponential-detrending fluctuations and the historical record (§3.1) exemplify the problems that can arise from inadequate trends.

In this paper we apply our method to only one variable. Many extensions are possible, such as applying our detrending method sector by sector, or to different countries. Because our analysis is univariate, the shocks we isolate are convolutions of demand and supply shocks. Because the relative growth trends of different variables are restricted along the balanced growth path, applying our methodology to a multivariate setting can be a fruitful avenue to identify supply and demand shocks.

Finally, our relaxation that output is log concave rather than log linear has a natural economic interpretation: the production function of technology can display diminishing marginal returns. Diminishing marginal returns are ubiquitous in economics; there is no reason to believe that technological progress should be exempt. A log linear balanced growth path is not an outcome of any model of endogenous or exogenous growth. It has always been an assumption imposed on the models because it was thought to be an empirical fact. Our results are part of a growing body of research that decisively overturns that fact.

References

Baxter, M., King, R.G., 1999. Measuring Business Cycles: Approximate Band-Pass Filters for Economic Time Series. The Review of Economics and Statistics 81, 575–593. doi:10.1162/003465399558454.

Bender, C.M., Orszag, S.A., 1999. Asymptotic Methods and Perturbation Theory. Advanced Mathematical Methods for Scientists and Engineers, Springer.

Blanchard, O., Fischer, S., 1989. Lectures on Macroeconomics. MIT Press.

Brillinger, D., 2001. Time Series: Data Analysis and Theory. Classics in Applied Mathematics, Society for Industrial and Applied Mathematics.
Brockwell, P.J., Davis, R., 2016. Introduction to Time Series and Forecasting, 3rd ed. Springer.
Burns, A., Mitchell, W.C., 1946. Measuring Business Cycles. National Bureau of Economic Research.
Cochrane, J.H., 2005. Time Series for Macroeconomics and Finance. URL: https://www.johnhcochrane.com/research/notes.
Comin, D., Gertler, M., 2006. Medium-term business cycles. American Economic Review 96, 523–551.
Elitzur, M., 2022. A general description of growth trends. Stats 5, 111–127.
Elitzur, M., Kaplan, S., Željko Ivezić, Zilberman, D., 2021. The impact of policy timing on the spread of covid-19. Infectious Disease Modelling 6, 942–954. doi:https://doi.org/10.1016/j.idm.2021.07.005.
Elitzur, M., Kaplan, S., Zilberman, D., 2020. Hindered growth. Journal of Economic Dynamics and Control 111, 103807. doi:https://doi.org/10.1016/j.jedc.2019.103807.
Eo, Y., Morley, J., 2022. Why Has the U.S. Economy Stagnated since the Great Recession? The Review of Economics and Statistics 104, 246–258. doi:10.1162/rest_a_00957.
Gabaix, X., 2011. The granular origins of aggregate fluctuations. Econometrica 79, 733–772. doi:https://doi.org/10.3982/ECTA8769.
Griliches, Z., 1957. Hybrid corn: An exploration in the economics of technological change. Econometrica 25, 501–522.
Groth, C., Koch, K., Steger, T., 2010. When economic growth is less than exponential. Econ Theory 44, 213–242.
Hamilton, J.D., 2018. Why You Should Never Use the Hodrick-Prescott Filter. The Review of Economics and Statistics 100, 831–843.
Hamilton, J.D., 2020. Time Series Analysis. Princeton University Press.
Harvey, A.C., Jaeger, A., 1993. Detrending, stylized facts and the business cycle. Journal of Applied Econometrics 8, 231–247. doi:https://doi.org/10.1002/jae.3950080302.
Hodrick, R.J., Prescott, E.C., 1981. Postwar U.S. Business Cycles: An Empirical Investigation. Technical Report. Northwestern University working paper.
Hodrick, R.J., Prescott, E.C., 1997. Postwar U.S. business cycles: An empirical investigation. Journal of Money, Credit and Banking 29, 1–16.
Ivezić, Ž., Connolly, A., Vanderplas, J., Gray, A., 2014. Statistics, Data Mining and Machine Learning in Astronomy. Princeton University Press.
Kingsland, S., 1982. The refractory model: The logistic curve and the history of population ecology. The Quarterly Review of Biology 57, 29–52.
Kocsis, T., Kovács-Székely, I., Anda, A., 2017. Comparison of parametric and non-parametric time-series analysis methods on a long-term meteorological data set. Central European Geology 60, 316–332. doi:10.1556/24.60.2017.011.
Mendoza, E.G., 1991. Real business cycles in a small open economy. The American Economic Review 81, 797–818.
Nakamura, E., Steinsson, J., Barro, R., Ursúa, J., 2013. Crises and recoveries in an empirical model of consumption disasters. American Economic Journal: Macroeconomics 5, 35–74. doi:10.1257/mac.5.3.35.
Nelson, C.R., Plosser, C.R., 1982. Trends and random walks in macroeconmic time series: Some evidence and implications. Journal of Monetary Economics 10, 139–162.

doi:https://doi.org/10.1016/0304-3932(82)90012-5.
Pearl, R., 1924. The curve of population growth. Proceedings of the American Philosophical Society 63, i–iv.
Perron, P., 1989. The great crash, the oil price shock, and the unit root hypothesis. Econometrica 57, 1361–1401. URL: http://www.jstor.org/stable/1913712.
Pianka, E.R., 1970. On r- and k-selection. The American Naturalist 104, 592–597.
Press, W.H., Teukolsky, S.A., Vetterling, W.T., Flannery, B.P., 2007. Numerical Recipes 3rd Edition: The Art of Scientific Computing. 3 ed., Cambridge University Press.
Preston, S.H., Heuveline, P., Guillot, M., 2000. Demography: measuring and modeling population processes. Wiley-Blackwell.
Ravn, M.O., Uhlig, H., 2002. On Adjusting the Hodrick-Prescott Filter for the Frequency of Observations. The Review of Economics and Statistics 84, 371–376. doi:10.1162/003465302317411604.
Razzak, W., 1997. The hodrick-prescott technique: A smoother versus a filter: An application to new zealand gdp. Economics Letters 57, 163–168. doi:https://doi.org/10.1016/S0165-1765(97)00178-X.
Ron, D., 1998. Guest editor's introduction. Machine Learning 30, 5–6.
Schacht, R.M., 1980. Two models of population growth. American Anthropologist 82, 782–798.
VanderPlas, J.T., 2018. Understanding the lomb–scargle periodogram. The Astrophysical Journal Supplement Series 236, 16.
Williamson, S.H., 2024. What Was the U.S. GDP Then? MeasuringWorth. URL: http://www.measuringworth.com/usgdp/.
Winsor, C.P., 1932. The gompertz curve as a growth curve. Proceedings of the National Academy of Sciences, USA 18, 1–8. doi:10.1073/pnas.18.1.1.
Wooldridge, J.M., 2012. Introductory Econometrics (5th ed). South-Western Cengage Learning.

Appendix A. Growth: Some Generalities

Appendix A.1. Growth Stability

Denote by Q_1 a random point on a growth trajectory where the growth rate is $g_1 = g(Q_1)$. Consider a small perturbation δQ so that $Q = Q_1 + \delta Q$ and $g = g_1 + \delta Q(dg/dQ)|_{Q=Q_1}$. Inserting in the equation of growth (eq. 2.1) and retaining only terms to 1st order in δQ shows that the ratio $\delta\rho = \delta Q/Q_1$ varies according to

$$\frac{d\delta\rho}{dt} = \alpha g_1\, \delta\rho, \qquad \text{where} \quad \alpha = \frac{d\ln g}{d\ln Q}\Big|_{Q=Q_1} \tag{A.1}$$

Accelerated growth has $dg/dQ > 0$, therefore $\alpha > 0$ and $\delta\rho$ grows exponentially. That is, the perturbation δQ diverges exponentially away from the underlying growth pattern — accelerated growth is inherently unstable. In addition, it is limited to finite durations (Elitzur, 2022). In contrast, decelerated growth ($dg/dQ < 0$, i.e., $\ln Q$ is concave downward) has $\alpha < 0$ and the perturbation decays exponentially relative to the underlying Q.

Appendix A.2. Logistic Growth

The logistic function is employed in the analysis of many phenomena, including population dynamics (Pearl, 1924; Kingsland, 1982) and diffusion of technology (Griliches, 1957). It can be written in the form of eq. 2.4 with the normalized profile

$$h_\ell(w) = \frac{2}{1 + e^{-w}} \tag{A.2}$$

(Schacht, 1980). Since $h_\ell(w \to \infty) = 2$, logistic growth is bounded by the carrying capacity $K = 2Q_\tau$. This follows also from the growth rate (equations 2.1 and A.2)

$$g = \frac{\gamma}{1 + e^{w}} = \gamma\left(1 - \frac{Q}{K}\right), \tag{A.3}$$

which vanishes at $Q = K$. As a function of t, g only vanishes as e^{-t} when $t \to \infty$, remaining finite at all finite t.

Logistic growth meets all the criteria in Definition 1, thus the logistic is a member of the ensemble of hindering functions. From eq. (2.5), the logistic growth rate has $p(q) = (1 - \frac{1}{2}Q/Q_\tau)^{-1}$, whose series expansion obeys equation 2.8 with $\lambda_k = 1/2^k$. Therefore the logistic function h_ℓ can be described by eq. 2.10 with this infinite set of expansion coefficients.

Appendix A.3. Two Epochs of Growth

The dichotomy between the unhindered and hindered phases of growth displayed by the sth functions (Figure 1) is common to all hindering functions. The growth-rate variation is controlled by the profile p (eq. 2.5), which is monotonically increasing with $p(0) = 1$. Since $p < 1$ when $w < 0$, it can be neglected to leading order, yielding the universal pattern of exponential growth. In contrast, $w > 1$ implies $p > 1$ and $g \simeq \gamma/p$; every p-profile sets its own growth trajectory, which can differ greatly from exponential. Based on the logistic, ecological models of population growth introduced the concept of r- and K-selection (Pianka, 1970; Schacht, 1980), the respective equivalents of the unhindered and hindered growth phases. This terminology reflected the notation for r as the maximal intrinsic rate of natural increase (γ here) and K the carrying capacity. In fact, two epochs of growth are inherent to all hindered growth, not unique to the logistic.

A simple example illustrates the two epochs: Consider the introduction of apple seeds to a newly discovered desolate island. Some seeds will sprout, apple trees will produce new seeds and the tree population will start growing exponentially at a rate determined by the island climate, ground fertility, etc. This

initial rate is γ (Definition 1), sometimes referred to as the natural growth rate (Schacht, 1980). The second epoch of growth is reached when the tree population increases so much that tree crowding becomes a significant factor, causing the growth rate to decline. Growth is impacted by crowding when $Q > Q_\tau$, where Q is the number of trees and Q_τ is the ratio of the island's area to the area required by a single tree to grow without hindering from others. The unconstrained growth of the first epoch is now hindered by competition for resources. Similar behavior is displayed in many other contexts, triggered by different factors and reflecting decreasing marginal growth rate.

Appendix A.4. The Hindering Functions Ensemble

A diminishing growth rate implies that $g(t)$ is a decreasing function of t. Therefore, when time is extended backward, g is *increasing*. The boundedness criterion (Definition 1.3, §2) ensures that this increase never becomes a blow-up — every member of the hindering functions ensemble is free of singularities in the entire upper half of the (t, Q) plane. However, this is not the case for some other descriptions of diminishing growth rates suggested in the literature. Groth et al. (2010) proposed a general description for log concave growth based on the putative expression $\dot{g} = -\beta g^2$, where the free parameter $\beta > 0$ is a damping coefficient. They recognized that the derived growth functions are applicable only beyond a certain initial time and cannot be extended to earlier epochs, where they blow up. This is also the case for log quadratic growth, which is based on the conjecture $\dot{g} = -\beta$ and is employed in many studies (e.g. Mendoza, 1991).

For insight into this behavior we note that the monotonicity criterion (Definition 1.1) ensures that the growth process can be considered in terms of either t or $Q(t)$ and thus can be envisioned as a trajectory in the 3-dimensional (t, Q, g) space: the parameter along the trajectory is the independent variable t, specifying the point where the quantity is $Q(t)$ and its growth rate $g(t)$. Figure A.12 shows trajectories for representative hindering functions, with coordinate along the t-axis corresponding to the variable w (eq. 2.3), the Q-axis to $Q(w)/Q_\tau$, i.e., $h(w)$ (eq. 2.10), and the g-axis to $g(w)/\gamma$. Because of the finite range of g (Definition 1 and subsequent discussion), the trajectories of all hindering functions are fully contained within the horizontal slab bounded by the planes $g = 0$ and $g = 1$. The panels below the 3-dimensional plot show the projections on the three planes. The left one, panel (a), shows the variation of Q with t, same as the top panel in Figure 1. The two other panels show two views of the variation of g. The mid panel, (b), shows g as a function of t, same as the mid panel in Figure 1, while the right one, panel (c), shows g as a function of Q.

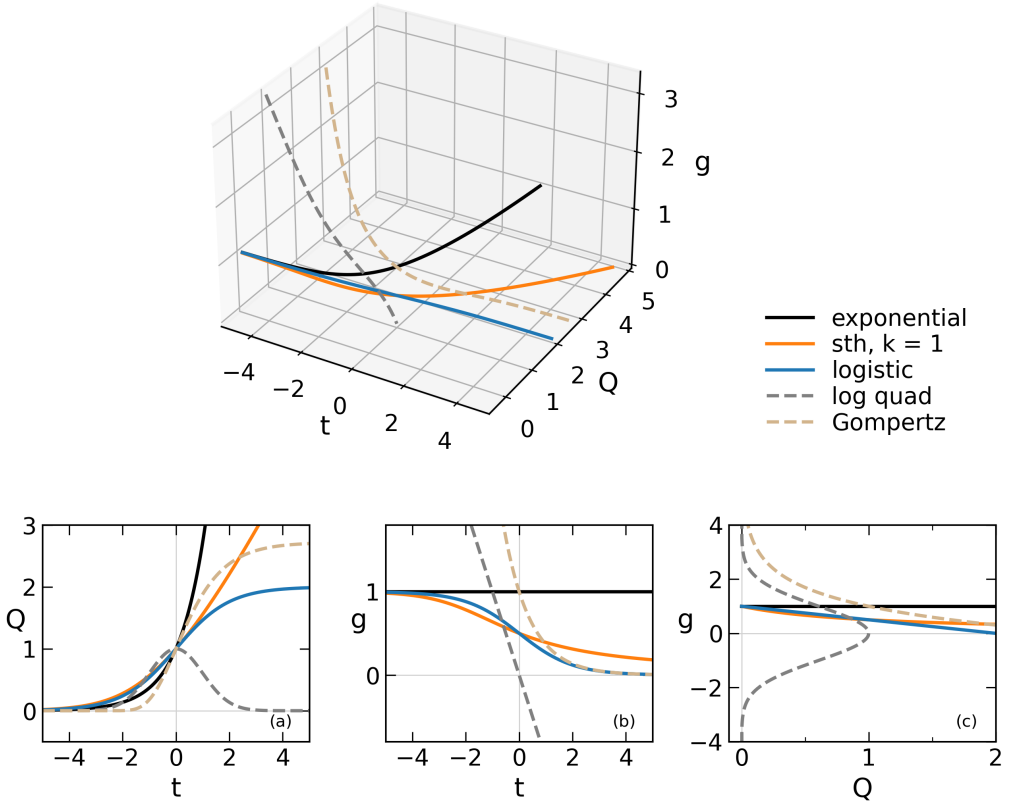


Figure A.12: *Top*: Growth processes evolve along trajectories in the 3-dimensional (t, Q, g) space. Projection on the t-Q plane (panel a) shows $Q(t)$, the growth pattern. Solid lines show three hindering functions: the exponential; sth, $k = 1$ (h_1, eq. 2.12); and logistic (h_ℓ, eq. A.2). All quantities in scaled units; see text for details. Projections on the two other planes show the growth-rate variation as a function of t (panel b) and of Q (panel c). Dashed lines show the log-quadratic (eq. A.4) and Gompertz function (eq. A.5), which are not members of the hindering functions ensemble. The growth rates of the hindering functions, and only the hindering functions, are singularity-free everywhere.

Of the three hindering functions plotted here, the exponential and $k = 1$ sth (linear hindering), were already shown in Figure 1 and their displayed variations in panels (a) and (b) replicate those in that figure. As the only function with a constant growth rate, the exponential trajectory stays in the horizontal plane $g = 1$. The ones for linear hindering and the logistic are descending toward the $g = 0$ plane. As functions of t, the growth rates of both vanish only in the limit $t \to \infty$, ensuring growth for all finite t. As a function of Q, the growth rate of

linear hindering vanishes only asymptotically as $1/Q$ while that of the logistic vanishes at a finite Q (= 2).

The hindering solution was derived by handling the growth rate as a function of $Q(t)$. We could instead utilize directly the independent variable t, assume $g(t)$ to be an analytic function and integrate its Taylor expansion to get a power series for $\ln Q$. The simplest case is linearly varying growth rate, $g = g_0 + \beta t$. Since growth stability requires $\beta < 0$ (Appendix A.1) we introduce $\sigma^2 = -1/\beta$, yielding the log quadratic growth $\ln Q = \ln Q_0 + g_0 t - \frac{1}{2}(t/\sigma)^2$. Growth ends at $t_f \equiv g_0\sigma^2$ where $g = 0$; at $t > t_f$, growth switches to contraction. The evolution of Q follows the Gaussian trajectory

$$Q = Q_f e^{-u^2/2}, \qquad g = -\frac{u}{\sigma} = \pm\frac{1}{\sigma}\left(-2\ln\frac{Q}{Q_f}\right)^{1/2} \tag{A.4}$$

with $u = (t - t_f)/\sigma$ and $\ln Q_f = \ln Q_0 + \frac{1}{2}(g_0\sigma)^2$. The switch from growth to contraction at $t = t_f$ is evident from the expression for g as a function of t; as a function of Q, the sign of g must be selected according to the branch of the growth pattern. Figure A.12 shows the growth trajectory of the log-quadratic, with u varying along the t-axis, Q/Q_f along the Q-axis and $g\sigma$ along the g-axis. The fundamental differences between the log-quadratic, shown with dashed line, and the hindering functions are on full display. Unlike the hindering functions, log-quadratic growth is confined to the $t < 0$ portion of the (t, Q, g) space; it is the only displayed growth pattern with limited duration. And while the trajectories of all hindering functions are fully contained within the horizontal slab bounded by the $g = 0$ and 1 planes, g of the log-quadratic not only becomes negative but also diverges linearly when $t \to \pm\infty$, and as $\pm(-\ln Q)^{1/2}$ when $Q \to 0$.

The singularities of log-quadratic growth afflict every model constructed from $\ln Q = P_N(t)$ with P_N a polynomial of order N. Whenever such a model produces stable decelerated growth, the growth rate diverges as $(-\ln Q)^{1-1/N}$ when $Q \to 0$; divergence is avoided only when $N = 1$, i.e., pure exponential growth, which is a member of the hindering-functions ensemble. The only other way to avoid divergence is to use an infinite t-series expansion instead of a finite polynomial; an example is the logistic, another member of the hindering-functions ensemble (Appendix A.2). However, an infinite t-expansion does not guaranty a singularity-free behavior. The Gompertz function, often used in demography and population studies (Preston et al., 2000), can be written as (Winsor, 1932)

$$Q = Q_0 e^{1-e^{-\nu}}, \qquad g = \frac{e^{-\nu}}{\theta} = \frac{1}{\theta}\left(1 - \ln\frac{Q}{Q_0}\right) \tag{A.5}$$

where $\upsilon = (t - t_0)/\theta$, with $\theta > 0$ a parameter, and $Q_0 = Q(t_0)$; the distribution's carrying capacity is $K = eQ_0$. Figure A.12 shows the Gompertz growth trajectory with υ varying along the t-axis, Q/Q_0 along the Q-axis and $g\theta$ along the g-axis. The growth rate diverges exponentially when $t \to -\infty$ and as $-\ln Q$ when $Q \to 0$.

The divergence problems are eliminated only when g obeys all the criteria in Definition 1. Growth functions that are not members of the ensemble of hindering functions are valid only in restricted regions of the (t, Q, g) space.

Appendix B. Technicalities

Appendix B.1. Fitting log data

It is common practice in macroeconomic studies to handle all data in log. In that case the fluctuations are defined as $X_t = \ln Y_t - \ln Q_t = -\ln Q_t/Y_t$ instead of eq. 3.1, where they are $1 - Q_t/Y_t$. The error estimate $\sum X_t^2$ is the same for the two approaches as long as $|X_t| \ll 1$ but differs for larger fluctuations.

We compared the two approaches and found that they generally produce similar results. However, the analysis in log yields significantly inferior results for the 20-year training set (§4.1). In particular, the overall error estimate $\sum \varepsilon_t^2$ of its predictions for the 2004–2023 data is 2.3 times larger than for the fit shown in Figure 8(a). This failure highlights the problems that can arise from the distortions introduced by the log transformation.

Appendix B.2. The Hindering Fitting Process

Our model, the set of coupled equations (3.1) and (3.2), is comprised of a deterministic trend, described by a hindering function, and stochastic fluctuations, described by an ARMA process. Fitting the data with every possible combination of hindering function and ARMA series is unrealistic. Instead, we adopt a sequential approach in the search for best-fitting model parameters.

A growth process is, naturally, driven by growth, therefore our first step is to determine the best-fitting hindering trend from equation (3.1). This is done with a systematic search of the ensemble of hindering functions, increasing the number of fitting parameters one by one—exponential, logistic, single-term hindering, two-term hindering, etc. For a given number of parameters, the best fit is the parameter set that minimizes the sum $\sum X_t^2$. This is the same as standard χ^2-minimization with the std of each point proportional to the data, $\sigma_t \propto Y_t$ (Press et al., 2007). It is important to note that we only seek the minimum of χ^2; its actual magnitude is irrelevant (no need to specify the proportionality constant in $\sigma_t \propto Y_t$). Comparing fits with different numbers of parameters requires a

separate test — since adding free parameters improves fitting (reduces χ^2) in itself, the statistical significance of such improvement must be evaluated. We use the Bayesian information criterion (BIC), which takes into account both the χ^2 and the number of free parameters of a fit (Ivezić et al., 2014; Brockwell and Davis, 2016). The best fitting trend is the one with the smallest BIC. In addition, fit improvement from added terms can also be assessed with another test. Comparing, for example, single- and two-term hindering, the former is a restricted form of the latter (the 2nd-term coefficient is restricted to zero), thus the statistical significance of the improvement from the 2nd term can be determined with the F-test (Wooldridge, 2012). We perform both tests and verify that their conclusions are in agreement.

Our search strategy gives precedence to simplicity of the trend function. For example, single-term hindering is fitted with $k = 1$, $k = 2$ and so on. The order k is increased as long as χ^2 keeps decreasing. The search is terminated once χ^2 starts increasing, identifying the smallest k to produce a local χ^2 minimum. Similarly, we stop adding hindering terms once the fit improvement becomes statistically insignificant.

In the second step we fit the fluctuations around the best-fitting trend with ARMA($\mathtt{p}, \mathtt{q}$) and search for the order ($\mathtt{p}, \mathtt{q}$) that minimizes the BIC of the fit. The search is performed over a square grid in the ($\mathtt{p}, \mathtt{q}$) plane, increasing its size until reaching a BIC minimum that falls inside the square. This search procedure identifies the point closest to the origin of the ($\mathtt{p}, \mathtt{q}$) plane around which the BIC of ARMA($\mathtt{p}, \mathtt{q}$) increases in every direction

In the final step we utilize the results from the first two steps to specify the model more narrowly. The trend Q_t is described by equation (2.4) with the hindering function h determined in the first step and γ, Q_τ and τ as free parameters. The ARMA process is specified by the order ($\mathtt{p}, \mathtt{q}$) determined in the second step, with the lag coefficients as free parameters. Minimization of $\sum \varepsilon_t^2$ determines all the free parameters simultaneously.

We perform all data analysis with software developed in-house employing the `python` programming language and tools from the packages `numpy`, `scipy` and `statsmodels`. The latter supplies all the ARMA analysis aids.

Appendix B.3. Spurious Hindering?

As noted in §3.1.1, purely exponential growth combined with different realizations of ARMA processes generated from a given set of parameters might display by chance a hindering trend. To estimate the probability of such spurious hindering we conducted Monte Carlo simulations, constructing synthetic data of fluctuations around exponential growth from the parameters of our best

fit to the US GDP (Table 1). In each simulation, fluctuations are generated from an ARMA process with the model ARMA parameters and Gaussian white noise whose width parameter is the residuals std. The resulting time series is inserted into eq. 3.1 for X_t together with an exponential with the constant growth rate γ used for Q_t. Solving the equation for Y_t produces our simulated process of exponential growth with fluctuations. The growth rate of this mock time series is computed by numerical differentiation and subjected to the MK test for a downward monotonic trend.

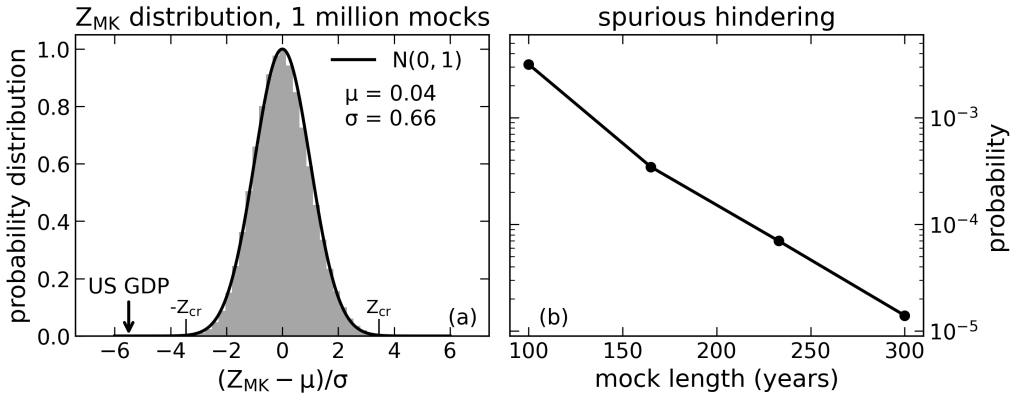


Figure B.13: Monte Carlo simulations of exponential growth with ARMA fluctuations defined by the parameters of our best fit to the US GDP. *Left*: Histogram of the MK statistic of the growth rates of 1 million mock datasets together with a normal distribution, both normalized to unity peak; the mean and std of the Z_{MK} distribution are listed. Z_{MK} of the US GDP growth rate and the critical values for detecting monotonic trends at 99% confidence level are marked; datasets to the left of $-Z_{cr}$ display hindered growth. *Right*: Variation with time-series length of the probability for spurious hindering that arises from fluctuations around an exponential trend.

The left panel of Figure B.13 shows the histogram of the MK statistic of 1 million simulated datasets together with markings of the critical Z for 99% confidence level, $Z_{cr} = 2.33$. Of the million mock datasets, 70 have $Z < -Z_{cr}$ and thus would pass as hindered growth processes. That is, with 99% confidence, the probability to display hindering by a growth process generated from fluctuations around an exponential trend with the parameters of our best fit is $7{\cdot}10^{-5}$. The most extreme MK statistic produced in the million simulations is -2.91. By contrast, the MK statistic of the US GDP growth rate is -3.60; the corresponding z-score is -5.50, expected once in 52 million simulations. The possibility that the US GDP growth rate is constant and only appears to be decreasing because of chance fluctuations can be dismissed as highly unlikely.

Spurious hindering arises when the fluctuations produce sufficiently long pe-

riods of growth-rate decline to affect the overall trend. The length of such coherent patterns is controlled by the duration of correlations between fluctuations, determined by the ARMA parameters. The probability for spurious hindering can thus be expected to vary with the ratio of this duration to the overall length of the time series — the longer the time series the smaller the ratio the lower the probability. The right panel of Figure B.13 shows this to be the case. It displays the variation of the probability for spurious hindering with length of the mock dataset when all other parameters are held fixed. The probability decreases exponentially from $3.2{\cdot}10^{-3}$ when the time series length is 100 years to $1.4{\cdot}10^{-5}$ at 300 years.

The simulation results provide guidance for our decomposition analysis. With the MK test proving a reliable indicator of long-term hindering, the detrending step progresses from the exponential to hindering functions only when the test detects growth deceleration. As shown in §2, a hindering function is guaranteed to fit any hindered growth process. However, such a fit need not imply a hindering trend in and of itself — a hindering function can provide the best fit even without long-term hindering if the fluctuations produce a sufficiently long downward trend toward the end of the time series. Our simulations produced ample evidence for such cases, and the MK test eliminated virtually all of from considerations for hindering. The 70 cases that did pass the test arose from rare coincidences among different segments of the mock data. The MK test is essential for separating chance fluctuations at the tail end of a time series from actual long-term hindering.

The simulations described here are the last step of our analysis method. Once the ARMA fit is determined, a Monte Carlo simulation is conducted as an a posteriori check to determine the probability of spurious hindering and verify the self-consistency of the decomposition.

Appendix B.4. Residual ACF and PSD

A decomposition is successful if its residuals ε_t are compatible with random noise. A noise process is not correlated with its own values at any lag other than 0, thus a time series is indistinguishable from noise when its acf vanishes for all lags ℓ other than 0. Testing for such δ-function shape presents two numerical challenges. First, because of the series finite size, the higher the lag the smaller the number of terms that enter into its ACF calculation. This leads to spurious structures, including waves, similar to problems afflicting Fourier analysis of finite bandwidth (VanderPlas, 2018). Such pitfalls are avoided by limiting ACF analysis to the first 25% of all lags (Brockwell and Davis, 2016, p 52). Second,

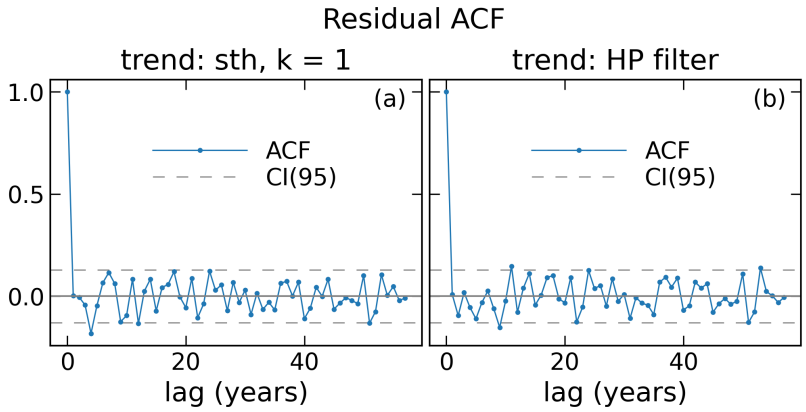


Figure B.14: Residual autocorrelation functions with lags to 58 for detrending methods as marked. At the 95% confidence level, the ACF of random noise are expected to fall within the indicated CI(95) boundaries.

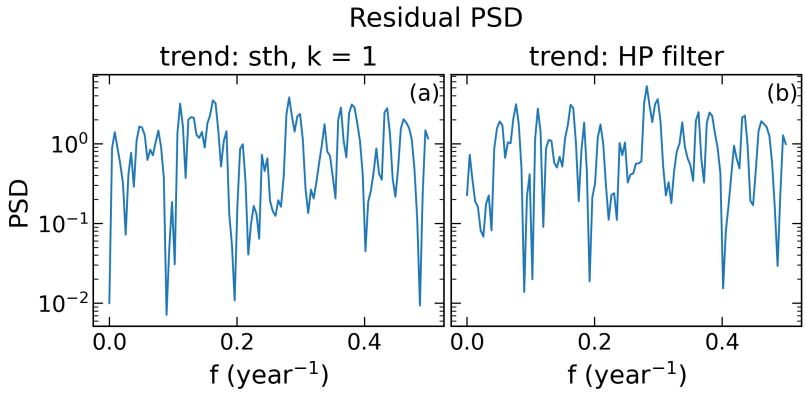


Figure B.15: Residual PSD (scaled units) for detrending methods as marked.

a zero value can be confirmed only within a confidence interval. At 95% confidence level, about 95% of random noise $\texttt{ACF}(\ell > 0)$ are expected to fall within the CI(95) boundaries, while roughly 5% could randomly protrude outside those boundaries (Brockwell and Davis, 2016). Our residual series have 233 members thus their ACF can be considered reliable for the first 58 lags. If no more than ~3 of those fall outside the CI(95) bounds, the residuals are consistent with random noise. Figure B.14 shows the residual autocorrelation functions for both hindering and HP decomposition methods. Both are compatible with random noise. Figure B.15 shows the PSD of each residual distribution. Both are compatible with the flat spectral density of white noise.