The Kent Equation: A Rate Law for the Onset of Local Biological Self-Instantiation

G Kent

June, 2026

Abstract

We introduce the Kent Equation, a model for the time at which a planetary civilization crosses the technical threshold required to fabricate a viable biological body, to specification, on its own surface. The model routes the crossing through a hazard driven by a substrate-maturity index $M(t)$—built from energy, general-purpose computation, and biological experimental throughput—raised to a responsiveness exponent and modulated by a conversion accelerant that measures how efficiently raw capacity is turned into solved developmental capability. Its primary object is not the crossing probability but its *rate of advance*: the logarithmic derivative of the hazard decomposes into a sum of measurable growth terms in which the unknown required substrate levels cancel exactly. From this single identity we derive four structural results: that the marginal return on conversion is gated by substrate maturity and is therefore back-loaded; that the percentage rate of approach and the absolute progress toward the threshold decouple; that the binding constraint is mobile and migrates with substrate maturity; and that the realized advantage of investing in capacity versus conversion is an elasticity-weighted race with an explicit crossover. We then show that recovering an actual crossing *date* requires committing to one quantity the rate law dispenses with—the gap between the binding input's current and required levels—and argue that this is the model's central use: it does not output a year, it localizes all timing disagreement to a single interpretable number and exhibits the sensitivity of the date to it.

1 Introduction

At what point in a civilization's technological development does it become possible to fabricate viable biological bodies, to specification, on the planet's own surface—to convert biological information into a living organism locally? We call a civilization that has crossed this threshold *instantiation-capable*, the threshold the *bootstrap point*, and we introduce the Kent Equation to model the time at which it is reached.

The equation is built around a distinction that recurs throughout. The *level* of progress —the probability that the threshold has been crossed by a given date—cannot be computed without knowing how mature each underlying capacity must become before the threshold is reached, and those required levels are deeply uncertain. The *rate* of progress—how fast the threshold is approaching—can be computed without them, because under the logarithmic derivative the required levels cancel. We therefore take the rate of advance as the model's primary object and recover the level, and any crossing date, only as a secondary construction that makes its single irreducible assumption explicit.

The model factors an intractable question into terms that different disciplines can estimate separately, in the tradition of factored estimators of the prevalence of life [1, 2, 3]. It differs from that tradition in form: it replaces a static product of probabilities with a time-indexed hazard, and its headline quantity is a velocity rather than a count.

We instantiate the model at human (multicellular vertebrate) scale and carry the organism class as a parameter elsewhere. Throughout we are explicit that this is a conditional timing model: it specifies what must be measured and how the measurements combine, not a calibrated forecast.

2 The model

2.1 The threshold and the crossing hazard

Definition 1 (Output)**.** Let the organism class Ω be a parameter of the model. $K_{\text{tech}}(t;\Omega)$ is the probability that a civilization has crossed the technical threshold required to fabricate, to specification and on its own surface, a viable biological body of class Ω, on or before time t, conditional on the substrate trajectory $M_{[0,t]}$:

$$K_{\text{tech}}(t;\Omega) = \mathbb{P}\big(\tau^* \le t \,\big|\, M_{[0,t]}\big)\,, \qquad K_{\text{tech}} \in [0,1], \tag{1}$$

where τ^* is the threshold-crossing time. We write $K_{\text{tech}}(t)$ for the human-scale instance.

Readiness is gated by the last of several necessary capabilities to mature, so the crossing time is a maximum over capability-crossing times. For each necessary capability c in a set $\mathcal{C}$, let τ_c be the first time c crosses the level required for instantiation; then

$$\tau^* = \max_{c\in\mathcal{C}} \tau_c. \tag{2}$$

The maximum is correct because the capabilities are necessary and conjunctive, and it natively represents the strong positive correlation among a single civilization's maturing capacities: every crossing time is routed through one shared driver. We model each crossing time through its hazard—the instantaneous rate of crossing given that it has not yet occurred—as a function of a common substrate-maturity trajectory $M(t)$ and a capability-specific conversion accelerant $A_c(t)$:

$$h_c(t) = A_c(t)\, M(t)^{\beta_c}, \qquad \mathbb{P}\big(\tau_c \le t \mid M_{[0,t]}\big) = 1 - \exp\left(-\int_{t_0}^{t} h_c(s)\, ds\right). \tag{3}$$

The exponent $\beta_c > 0$ encodes how responsive capability c is to substrate: a large β_c describes a threshold-like capability whose hazard stays near zero until substrate is well matured.

At human scale the binding capability is developmental viability—growing a specified genome into a supported, viable organism. The empirical record places it last: genome *writing* has reached single-celled-eukaryote scale [4, 5], advancing along the throughput-and-cost axis that capacity and machine inference accelerate, whereas developmental *viability* reaches only the two ends of gestation—ex utero embryo culture for a few days [6], and artificial-womb support confined to the final gestational third [7]—with continuous instantiation of a specified body absent. Development is gated less by computation than by empirical biological knowledge obtainable only by running development, so it has the higher exponent and the later crossing. We therefore take the canonical model to be the single binding gate, write $h(t) = A(t)\, M(t)^{\beta}$ for it (dropping the developmental subscript), and treat the earlier-crossing capabilities as necessary but non-binding prerequisites.

2.2 The substrate-maturity index

The driver must be measurable. We build it from three components, each chosen for a named causal pathway to developmental viability rather than a merely correlated upward trend: energy E (generation capacity and falling real cost, upstream of all computation and experiment); general-purpose computation C (cost per operation and installed capacity, which carries the inference that performs sequence design and developmental simulation); and biological experimental throughput B (automated embryo culture, organoids, imaging, closed-loop wet-lab automation, which generates the developmental knowledge viability requires).

The components are combined by a weighted geometric mean, and each is normalized not to an arbitrary present-day maximum but to the level x_i^* it must reach for the gate:

$$M(t) = \prod_{i \in \{E,C,B\}} \left(\frac{x_i(t)}{x_i^*} \right)^{w_i}, \qquad \frac{x_i(t)}{x_i^*} \leq 1 \text{ (capped per component)}, \qquad w_i > 0. \tag{4}$$

Two choices are substantive. The geometric mean treats the components as *complementary* rather than substitutable: all layers are required, and a near-absence of any one drives the index toward zero. A civilization with abundant energy and experimental throughput but negligible computation cannot perform the design and developmental inference on which fabrication depends, and the geometric mean correctly assigns it near-zero maturity, where a weighted sum would wrongly score it as moderately mature. Normalizing to the required levels relocates the model's principal free assumption from an arbitrary scaling maximum to the x_i^*—a substantive scientific question on which domain experts can disagree productively. We do not claim the required levels are known; we claim only that this is the right place for the uncertainty to live, and Section 2.5 shows the rate of approach does not depend on them at all.

2.3 The conversion accelerant

The components of $M(t)$ are capacity variables: they measure how much raw substrate exists, not how effectively it is converted into solved problems. Conversion is a distinct capability, and it is the natural locus of machine intelligence. We model it as an accelerant on the hazard rather than a component of capacity.

Proposition 1 (Two-sided accelerant)**.** *The accelerant satisfies $A(t) > 0$, with $A = 1$ the no-net-acceleration baseline at which substrate is converted at its brute-force rate. Values $A > 1$ represent net acceleration; values $0 < A < 1$ represent net drag—institutional, epistemic, or safety friction that converts substrate less efficiently than baseline. No floor $A \geq 1$ is imposed.*

Remark 1 (Drag slows but, under a divergence condition, does not block)*.* Because the baseline hazard is positive for mature substrate, even an epoch of net drag still crosses the threshold provided $\int^{\infty} A(s)\, M(s)^{\beta}\, ds = \infty$. Acceleration and drag set the *speed* of crossing; the divergence of the cumulative hazard, not the mere positivity of A, is what guarantees the crossing occurs. A sufficiently fast-decaying accelerant is the one regime in which crossing is not guaranteed.

2.4 The level form

Collecting the single binding gate, the cumulative (level) form of the Kent Equation is the crossing-time distribution

$$K_{\text{tech}}(t) = 1 - \exp\left(- \int_{t_0}^{t} A(s)\, M(s)^{\beta}\, ds \right). \tag{5}$$

It is non-decreasing, departs from zero at the epoch t_0, and rises toward one as substrate matures. Equation (5) is the object one would evaluate to read a crossing date, and it cannot be evaluated without the required levels x_i^* that define M through (4). This is the motivation for taking its derivative.

2.5 The rate law

The model's headline object is the logarithmic derivative of the hazard—the rate at which the instantiation clock accelerates. Taking $\ln h = \ln A + \beta \ln M$ and differentiating in time, with $\ln M = \sum_i w_i \big(\ln x_i - \ln x_i^*\big)$, the constants $\ln x_i^*$ vanish:

$$\boxed{\frac{\dot{h}}{h} = \frac{\dot{A}}{A} + \beta \sum_{i \in \{E,C,B\}} w_i \frac{\dot{x}_i}{x_i}} \tag{6}$$

Writing $r_A = \dot{A}/A$ for the conversion growth rate and $r_i = \dot{x}_i/x_i$ for the substrate component growth rates, (6) reads $\dot{h}/h = r_A + \beta \sum_i w_i r_i$. This is the Kent Equation in the form we use throughout. Its terms are the growth rates of quantities with real, long-run time series, and—unlike the level (5)—it contains no required level. *How fast the threshold is approaching can be measured without knowing how far away it is.*

3 Insights

The remaining results are consequences of (6) and of the multiplicative substrate-conversion structure of the hazard. None requires calibration.

3.1 A disciplinary decomposition

The four kinds of term in (6) route to four different bodies of expertise (Table 1). Three of the four—the substrate growth rates r_i, given the weights—are directly measurable, and the fourth, the responsiveness β, is a single constant rather than a function. The decomposition does the work a factored estimator should: it isolates sub-questions and localizes disagreement to the term that carries it.

Term	Meaning	Home discipline	Status
$r_A = \dot{A}/A$	conversion-efficiency growth	machine intelligence	estimable
β	developmental responsiveness	developmental biology	single constant
w_i	relative substrate load	cross-disciplinary	weights
$r_i = \dot{x}_i/x_i$	substrate growth (energy, compute, throughput)	energy, industry, econ.	measurable

Table 1: The terms of the rate law (6) and where each is estimated.

3.2 The return on conversion is substrate-gated

The hazard's sensitivities to its two inputs are

$$\frac{\partial \ln h}{\partial \ln A} = 1, \qquad \frac{\partial \ln h}{\partial \ln M} = \beta, \qquad \frac{\partial h}{\partial A} = M^{\beta}. \tag{7}$$

The absolute marginal return on conversion is M^β. For a threshold-like developmental capability, $\beta > 1$, so M^β is convex and stays near zero until M is close to one (Figure 1). The consequence is sharp: intelligence applied to body-growing buys almost nothing while the wet-lab substrate is immature, and a great deal once it has matured. The payoff to conversion is back-loaded, and the size of the back-loading is governed by β.

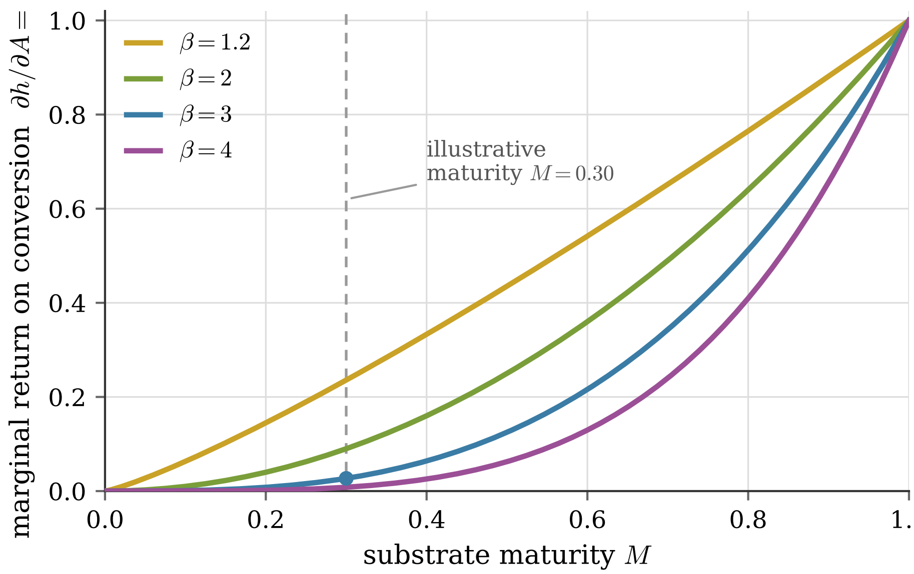

Figure 1: The return on conversion $\partial h/\partial A = M^\beta$ as a function of substrate maturity M, for several responsiveness exponents. For $\beta > 1$ the curve is convex and near zero across most of the range: conversion gains move the hazard appreciably only once substrate is mature. The values are illustrative.

3.3 Rate and level decouple

A racing component lifts the rate of approach without lifting the absolute progress toward the threshold. The mechanism is the complementarity of the index.

Lemma 1 (The least-mature component caps maturity)**.** *With capped ratios* $\rho_i = \min(x_i/x_i^*, 1) \le 1$ *and* $M = \prod_i \rho_i^{w_i}$,

$$M \le \min_j \rho_j^{w_j} \le \min_j \rho_j. \tag{8}$$

Proof. Every factor $\rho_i^{w_i} \le 1$, so for each fixed j, $M = \rho_j^{w_j} \prod_{i \ne j} \rho_i^{w_i} \le \rho_j^{w_j}$. Taking the minimum over j gives the bound; $\rho_j^{w_j} \le \rho_j$ follows from $\rho_j \le 1$. $\square$

A single lagging component therefore bounds M, and with it the absolute hazard AM^β, near zero regardless of how fast the other components grow. Yet that lagging component is irrelevant to which term *dominates the rate*: a fast-growing component contributes $\beta w_i r_i$ to $\dot{h}/h$ in (6) whatever the level of M. Hence the percentage acceleration can be large, and dominated by a racing input, while the threshold remains as distant in absolute terms as the laggard allows. The rate is scale-free; the level is not. This is the precise sense in which co-movement of indices is nearly uninformative about proximity: empirical content lives in the timing and ordering of crossings, not in shared direction of growth, and an observer tracking the fastest-moving input will systematically misjudge how close the threshold is.

3.4 Two levers and their crossover

Rewriting (6) as $\dot{h}/h = r_A + \beta\,(\dot{M}/M)$ exposes two levers on the rate of approach. The conversion lever r_A has unit elasticity and is algorithmic and fast; the substrate lever $\beta\,(\dot{M}/M)$ has elasticity $\beta > 1$ but is capital-bound and slow. Their realized contributions are elasticity times growth, and they trade places at the crossover

$$r_A = \beta\,\frac{\dot{M}}{M} = \beta \sum_i w_i r_i. \tag{9}$$

Below the crossover the substrate lever dominates the rate of approach; above it, conversion does. The decision content is immediate: the higher-leverage investment shifts from capacity to conversion as the crossover is passed, and (9) states exactly where the passage occurs in terms of observable growth rates.

3.5 The binding constraint is mobile

The identity of the binding capability is not fixed; it is a function of substrate maturity and relative acceleration. For two capabilities—development and an earlier-crossing prerequisite, synthesis—the ratio of hazards is

$$\frac{h_{\text{dev}}}{h_{\text{syn}}} = \frac{A_{\text{dev}}}{A_{\text{syn}}}\, M^{\,\beta_{\text{dev}}-\beta_{\text{syn}}}. \tag{10}$$

A lower hazard means a later crossing, hence the binding gate. With $\beta_{\text{dev}} > \beta_{\text{syn}}$ and $M < 1$, development binds at comparable acceleration. It ceases to bind—the constraint migrates to synthesis—once

$$\frac{A_{\text{dev}}}{A_{\text{syn}}} > M^{-(\beta_{\text{dev}}-\beta_{\text{syn}})}, \qquad \text{equivalently at maturity} \qquad M_{\text{flip}} = \left(\frac{A_{\text{dev}}}{A_{\text{syn}}}\right)^{-1/(\beta_{\text{dev}}-\beta_{\text{syn}})}. \tag{11}$$

"What to watch" is therefore conditional. An observer who treats the binding gate as permanent will misjudge proximity precisely once investment begins to move the bottleneck, and (10) quantifies how hard the binding capability would have to be accelerated, relative to the prerequisite, for the migration to occur.

4 From rate to time

The rate of approach is measurable; recovering a crossing *date* is not free. Because the binding hazard is sharply increasing (β large), the level (5) rises steeply near the time M approaches one, which—by the capped geometric mean (4)—occurs when the last component reaches its required level. Under exponential trajectories $x_i(t) = x_i(t_0)\,e^{r_i(t-t_0)}$, component i reaches x_i^* at $t_0 + \ln(x_i^*/x_i(t_0))/r_i$, so

$$\tau^* \;\approx\; t_0 + \max_i \frac{\ln\!\big(x_i^*/x_i(t_0)\big)}{r_i}. \tag{12}$$

Each input contributes a time equal to its gap-to-requirement in natural-log units divided by its growth rate—equivalently, the number of doublings it must still close times its doubling time. The binding input *in time* is the one maximizing this ratio, which need not be the slowest-growing input: a fast-growing input with a large gap can bind. The accelerant pulls τ^* modestly earlier, but for large β the crossing is dominated by the time at which $M \to 1$, so the date is governed by the gap-and-growth structure of (12) with A as a tail correction.

Equation (12) isolates where the uncertainty in a date actually sits. The growth rates r_i are measurable; therefore all timing uncertainty resides in the gaps $\ln(x_i^*/x_i(t_0))$, and in practice in the single gap on the binding input. This is the model's central methodological claim.

Proposition 2 (The deliverable is a localized disagreement, not a date)**.** *The model does not output a year. It converts the question "when is the threshold crossed?" into "how many doublings short is the binding capability, and how fast is it growing?", and it exhibits the sensitivity of the date to that one quantity. Two analysts who agree on the growth rates but disagree on the date must, by* (12)*, disagree about a gap—a single interpretable number per input, dominated by the one binding input's gap—rather than arguing past one another about timelines. The model's product is a disciplined account of why a given date follows and what one belief would have to change to move it.*

Figure 2 makes Proposition 2 concrete. For plausible growth rates of the binding input, the crossing date is a transparent function of the assumed gap: fix the gap and the date band is fixed; change the gap and the entire fan of dates shifts. The figure is not a forecast but a sensitivity surface—the object the model is built to produce.

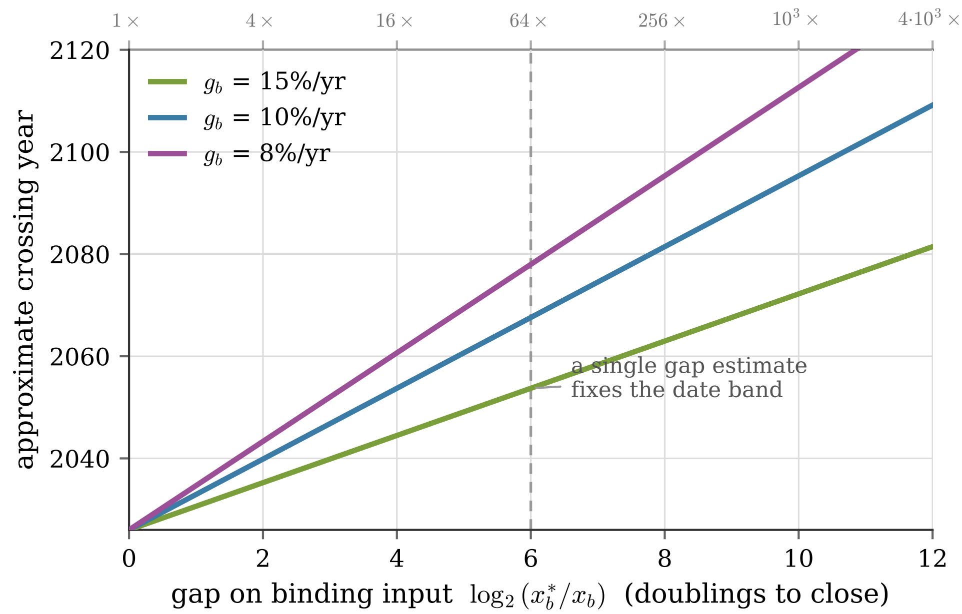

Figure 2: Crossing date from (12) as a function of the gap on the binding input (in doublings to close), for several growth rates, anchored to the present. The date is a function of one assumption: a single gap estimate fixes the date band, and the spread across plausible gaps dominates the spread across growth rates. Top axis shows the gap as a multiplicative factor. Values are illustrative.

5 Uses

Dated, falsifiable sub-predictions. The model commits to claims testable well before the terminal crossing. The sign and curvature of $\dot{h}/h$ are observable from the public substrate series; the ordering claim—that the prerequisite capability crosses earlier in time than the binding one—is checkable; and the conjunctive structure predicts a late-inflecting, S-shaped trajectory in developmental indicators rather than the linear extrapolation a naive reading of current culture durations would suggest. Any of these can fail, which is the sense in which the model carries empirical content.

Indicator selection. With the structure fixed, (12) identifies which measurement most tightens an estimate of the crossing date: the gap on the binding input. It directs monitoring to developmental-viability indicators—the duration over which embryos can be cultured ex utero, the completeness of synthetic embryo models, and the gestational coverage of artificial-womb technology—rather than to genome-synthesis milestones, which by Section 2.5 and (10) crosses earlier and is not the binding signal.

Localizing disagreement. By Proposition 2, forecasters who differ on the date can locate their disagreement in a single parameter and argue about it directly. This is the use we regard as primary, and the reason the model earns its structure despite producing no calibrated number.

6 Limitations

- **The gap is the load-bearing judgment.** Equation (12) relocates the timing uncertainty to the gap on the binding input but does not remove it; for development the gap is genuinely multi-dimensional (culture duration is only one axis alongside morphogenesis, immune integration, and viability confirmation for a novel organism), and reducing it to a scalar is itself a modelling choice.
- **Co-trending is mitigated, not eliminated.** In an era of broad technological growth nearly every index rises together, so correlation between M and the gated capability is nearly free; the model's empirical content lives in its timing and ordering claims, a real constraint on testability.
- **The responsiveness β is unknown.** The power-law form makes the responsiveness a single constant rather than a state-dependent elasticity, but estimating it is an open problem in developmental biology.
- **Crossing certainty requires a divergence condition.** The guarantee that even a drag regime crosses depends on $\int^{\infty} AM^{\beta}\,ds = \infty$, not on the positivity of A alone.
- **The model is a structure, not a calibration.** The contribution is the equation and its decomposition; any specific quantities substituted, including those in the figures, are illustrative.

7 Conclusion

The Kent Equation is a rate law for the onset of local biological self-instantiation. Its measurable object is the velocity of approach (6), in which the unknown required substrate levels cancel; its structural results—the substrate-gated, back-loaded return on conversion, the decoupling of rate from level, the mobile binding constraint, and the elasticity-weighted lever crossover—follow from the multiplicative substrate-conversion form of the hazard. Recovering a crossing date requires committing to one quantity the rate law dispenses with, the gap on the binding input, and the model's central use is exactly this: not to announce a year, but to localize all timing disagreement to a single interpretable number and to display how the date depends on it. We have specified what must be measured and where each measurement enters, and leave the empirical estimation of the gap, the growth rates, and the responsiveness to subsequent work.

References

[1] F. D. Drake, "The radio search for intelligent extraterrestrial life," in G. Mamikunian and M. H. Briggs (eds.), *Current Aspects of Exobiology*, Pergamon Press, 1965.

[2] C. Sagan and I. S. Shklovskii, *Intelligent Life in the Universe*, Holden-Day, 1966.

[3] D. A. Vakoch and M. F. Dowd (eds.), *The Drake Equation*, Cambridge University Press, 2015.

[4] D. G. Gibson *et al.*, "Creation of a bacterial cell controlled by a chemically synthesized genome," *Science* **329** (2010) 52–56.

[5] N. Annaluru *et al.*, "Total synthesis of a functional designer eukaryotic chromosome," *Science* **344** (2014) 55–58.

[6] A. Aguilera-Castejón *et al.*, "Ex utero mouse embryogenesis from pre-gastrulation to late organogenesis," *Nature* **593** (2021) 119–124.

[7] E. A. Partridge *et al.*, "An extra-uterine system to physiologically support the extreme premature lamb," *Nature Communications* **8** (2017) 15112.