A Formal Economic Model of Technological Sovereignty

Analytical Appendix to "Two Kinds of Sovereignty"

Abstract

This document presents a three-layer formal economic model that unifies the core mechanisms of the essay "Two Kinds of Sovereignty": the exploitation/exploration allocation tradeoff (March, 1991), path-dependent lock-in (Arthur, 1989), induced innovation under constraint (Hicks, 1932; Hayami & Ruttan, 1971), and techno-economic paradigm windows (Perez, 2002). Layer 1 solves the social planner's optimal control problem analytically. Layer 2 validates the analytical results through agent-based evolutionary simulation. Layer 3 translates both into actionable policy instruments. The model is calibrated to European data and produces three principal outputs: optimal allocation ratios between present and future sovereignty, threshold conditions for when designed constraint dominates laissez-faire, and the shadow price of delay in contested technological domains.

Contents

1 Layer 1: Analytical Core

1.1 Setup

A social planner allocates a unit resource flow R between two activities at each moment t:

- **Exploitation** αR: secures and domesticates existing technology. Produces present sovereignty capital K_p.
- **Exploration** $(1-\alpha)R$: invests in new technological paradigms. Produces future sovereignty capital K_f.

1.2 State Variables

The system is characterised by four state variables:

Variable	Description	Range
$K_p(t)$	Present sovereignty stock	≥ 0
$K_f(t)$	Future sovereignty stock	≥ 0
$D(t)$	Dependency level	$[0, 1]$
$W(t)$	Paradigm window openness	$[0, 1]$

1.3 Dynamics

Definition 1 (Present Sovereignty Accumulation)**.**

$$\frac{dK_p}{dt} = A \cdot (\alpha R)^{\kappa} - \delta_p K_p - \varphi(D) K_p \tag{1}$$

where A is total factor productivity for exploitation, $\kappa \in (0,1)$ imposes diminishing returns, δ_p is the base depreciation rate, and $\varphi(D) = \varphi_0 D^2$ is dependency-accelerated depreciation (convex in D).

Definition 2 (Future Sovereignty Accumulation)**.**

$$\frac{dK_f}{dt} = W(t) \cdot [(1-\alpha)R]^{\beta} \cdot h(\sigma) - \delta_f K_f \tag{2}$$

where $\beta \in (0,1)$ governs exploration returns, and $W(t)$ is the Perez window function. Crucially, exploration investment has zero marginal return once the window closes. For analytical tractability, Layer 1 treats the economy as targeting a single paradigm window at a time; the multi-domain portfolio is handled in Layer 3.

Definition 3 (Dependency Dynamics (Arthur Increasing Returns))**.**

$$\frac{dD}{dt} = \lambda(\bar{D} - D)D - \eta K_p D \tag{3}$$

This is a logistic process toward the lock-in ceiling $\bar{D}$, slowed by present sovereignty investment. The multiplicative D term in $\eta K_p D$ ensures $D \in [0,1]$: when $D \to 0$, the reduction term vanishes, preventing negative dependency.

Definition 4 (Perez Window Dynamics)**.**

$$W_j(t) = W_{0,j} \cdot \exp\left(-\gamma_j \cdot \max(0,\, t - t_{open,j})\right) \quad \text{for } t \geq t_{open,j} \tag{4}$$

Windows open at $t_{open,j}$ with initial openness $W_{0,j}$ and close exponentially at rate γ_j.

1.4 The Constraint Instrument σ

The essay's "virtual export restriction" enters the exploration production function as a non-monotonic modifier:

$$h(\sigma) = 1 + a\sigma \exp(-b\sigma), \quad a, b > 0 \tag{5}$$

Proposition 1 (Hump Shape of $h(\sigma)$)**.** *The function $h(\sigma)$ defined in* (5) *satisfies:*

1. $h(0) = 1$ *(no constraint, baseline productivity)*
2. *h is maximised at $\sigma^* = 1/b$ with $h(\sigma^*) = 1 + a/(be)$*
3. *$h(\sigma) \to 1$ as $\sigma \to \infty$*
4. *$h(\sigma) > 1$ for all $\sigma > 0$*

Proof. $h'(\sigma) = a\exp(-b\sigma)(1 - b\sigma)$. Setting $h'(\sigma) = 0$ gives $\sigma^* = 1/b$. $h''(\sigma^*) = -ab\exp(-1) < 0$, confirming a maximum. For large σ, $a\sigma\exp(-b\sigma) \to 0$, so $h \to 1$. Since $a\sigma\exp(-b\sigma) > 0$ for all $\sigma > 0$, we have $h > 1$ everywhere. □

This hump shape formalises the essay's three selection pressure lanes:

- **Absent pressure** ($\sigma = 0$): Baseline productivity, no induced innovation.
- **Designed pressure** ($\sigma \approx \sigma^*$): Constraint maximises exploration productivity.
- **Destructive pressure** ($\sigma \gg \sigma^*$): Excessive constraint degrades toward baseline.

1.5 The Planner's Problem

$$\max_{\alpha(t), \sigma(t)} \int_0^\infty e^{-\rho t} \cdot V(K_p, K_f, D)\, dt \tag{6}$$

subject to (1)–(3), where the flow value is:

$$V(K_p, K_f, D) = \omega K_p + (1 - \omega)K_f - \psi D^2 \tag{7}$$

Remark 1 (Rationale for the Linear-Quadratic Form). Sovereignty benefits are linear because each unit of K_p or K_f provides roughly constant marginal strategic value. Dependency costs are quadratic because the damage is convex: going from 50% to 60% dependent is substantially worse than 20% to 30%, reflecting non-linear lock-in and narrowing exit options. This form ensures the optimal policy is smooth (interior solution) rather than bang-bang.

1.6 Analytical Outputs

The model is solved numerically via forward simulation over a grid of (α, σ) values, with scalar optimisation refining the peak. Key outputs:

Proposition 2 (Optimal Allocation)**.** *Under the baseline calibration ($D_0 = 0.67$, $R = 0.022$, $\gamma = 0.1$), the optimal fixed allocation is $\alpha^* \approx 0.35$, with optimal constraint intensity $\sigma^* = 2.0$. This implies roughly 65% of sovereignty-oriented resources should flow to exploration (future sovereignty) rather than exploitation (present sovereignty).*

Proposition 3 (Shadow Price of Delay)**.** *The cost of delaying exploration investment by Δt years is:*

$$C(\Delta t) = V(\alpha^*, \sigma^*, W_0) - V(\alpha^*, \sigma^*, W_0 \cdot e^{-\gamma \Delta t}) \tag{8}$$

This cost is non-linear and accelerating: each additional year of delay costs more than the last, because windows close exponentially. Under baseline calibration, a 5-year delay costs approximately 3 times as much as a 1-year delay.

1.7 Regime Comparison

Three policy regimes produce strikingly different trajectories:

Regime	α	σ	V (discounted)	Interpretation
Status Quo (Presentism)	0.80	0.0	-2.55	Exploit-heavy, no constraint
Designed Pressure	0.35	2.0	3.93	Optimal allocation + constraint
Overreaction	0.20	15.0	-5.09	Excessive constraint

The designed pressure regime dominates both alternatives, yielding positive total value where both alternatives are negative. The status quo's negative value reflects the compounding cost of dependency under path-dependent lock-in.

2 Layer 2: Evolutionary Simulation

2.1 Purpose

The analytical core assumes a representative agent and smooth production functions. Layer 2 relaxes both, populating the economy with N heterogeneous firms to capture emergent dynamics: why some constraint regimes produce compounding capability while others fragment the innovation system.

2.2 Agent Structure

Each of $N = 500$ firms is characterised by:

Attribute	Description	Initial Distribution
$c_i(t)$	Capability stock	Log-normal, mean-normalised to 1
$r_i(t)$	Research orientation $[0,1]$	Beta$(2,5)$, exploitation-skewed
$x_i(t)$	External dependency ratio	Beta$(5,2)$, high-dependency-skewed
$s_i(t)$	Market share	Uniform: $1/N$

2.3 Innovation Process

Each period ($\Delta t = 1$ year), firms draw from two innovation lotteries:

Exploitation. Probability $p_0 \cdot c_i^\nu$, payoff $\varepsilon_{\text{exploit}} \cdot c_i \cdot x_i \cdot (1 - \sigma x_i)$. High probability, low variance. The constraint factor $(1 - \sigma x_i)$ means highly dependent firms face an immediate productivity shock under designed pressure.

Exploration. Probability $q_0 \cdot c_i^\nu \cdot W(t)$, payoff drawn from Pareto$(\xi) \cdot \varepsilon_{\text{explore}} \cdot c_i \cdot (1 - x_i)$. Low probability, fat-tailed payoff. Exploration success also reduces the firm's dependency ratio.

2.4 Selection Mechanism

Market shares update each period via:

$$\tilde{s}_i(t+1) = s_i(t) \cdot \left(\frac{c_i(t)}{\bar{c}(t)}\right)^\theta, \quad s_i(t+1) = \frac{\tilde{s}_i(t+1)}{\sum_j \tilde{s}_j(t+1)} \tag{9}$$

Firms with $s_i < s_{\min}$ exit and are replaced by new entrants drawn from the initial distributions.

2.5 Emergent Regimes

Under different σ values, three qualitatively distinct system dynamics emerge, corresponding to the essay's three selection pressure lanes:

1. $\sigma = 0$ **(Absent Pressure):** Firms drift toward high dependency. Aggregate capability grows slowly and undirected. Efficient but fragile.
2. $\sigma \approx \sigma^*$ **(Designed Pressure):** High-capability firms redirect toward exploration. Aggregate capability dips initially then compounds as alternative pathways mature.
3. $\sigma \gg \sigma^*$ **(Destructive Pressure):** Most firms lack the absorptive capacity to redirect. Mass substitution with inferior alternatives. Aggregate capability degrades.

2.6 Cross-Layer Validation

Monte Carlo sweeps across σ values confirm the hump-shaped relationship between constraint intensity and aggregate terminal capability, validating the analytical $h(\sigma)$ assumption. The simulation also reveals distributional dynamics invisible to Layer 1: the *capability threshold* below which designed pressure destroys rather than redirects firm-level innovation.

Proposition 4 (Capability Threshold)**.** *There exists a minimum initial capability $c_{\min}(\sigma)$, increasing in σ, below which firms lose capability under designed pressure relative to the unconstrained baseline. This threshold determines the absorptive capacity precondition for successful strategic necessity programmes.*

3 Layer 3: Policy Design Interface

3.1 Observable State Variables

The planner does not observe K_f or $W(t)$ directly. The policy interface maps observable proxies to recommendations:

Observable	Proxy for	Empirical Anchor
D_{obs}	Dependency $D(t)$	Cloud market share: 65–70%
C_{obs}	Capability distribution	R&D intensity: 2.2% GDP
W_{obs}	Window openness	Maturity curves, VC acceleration

3.2 Policy Instruments

Instrument 1: Allocation Ratio α^*. How much sovereignty spending goes to exploitation vs. exploration. Given $(D_{\text{obs}}, C_{\text{obs}}, W_{\text{obs}})$, the model outputs α^* and its sensitivity to measurement error.

Instrument 2: Constraint Intensity σ^*. The "virtual export restriction" strength, scaled by observed capability: $\sigma^* = \sigma^*_{\text{base}} \cdot C_{\text{obs}}$. Low-capability systems receive weaker designed pressure to avoid the destructive regime.

Instrument 3: Domain Selection. Candidate paradigm windows scored on:

$$S_j = w_W \cdot W_{j,\text{remaining}} + w_C \cdot C_j + w_S \cdot V_j^{\text{strategic}} \tag{10}$$

with weights $(w_W, w_C, w_S) = (0.4, 0.3, 0.3)$.

3.3 Domain Assessments

Domain	$W_{\text{remaining}}$	C_j	$V_j^{\text{strategic}}$	Score
Fusion Energy	0.8	0.5	0.9	0.74
Scientific AI	0.7	0.4	0.8	0.64
Quantum	0.6	0.6	0.7	0.63

Fusion scores highest due to a newly opened window and very high strategic value. Scientific AI and quantum are closely ranked, with different profiles: quantum has stronger existing European capability, while scientific AI has a wider remaining window.

4 Calibration

Parameter	Value	Source	Layer
D_0	0.67	Synergy Research Group	1, 2, 3
R	0.022	European Investment Bank	1
κ	0.6	Cobb-Douglas convention	1
β	0.5	Exploration returns	1
δ_p	0.05	Annual depreciation	1
δ_f	0.03	Knowledge decay	1
$\bar{D}$	0.95	Lock-in ceiling	1
λ	0.3	Lock-in speed	1
a	2.0	Hump amplitude	1
b	0.5	Hump decay $\Rightarrow \sigma^* = 2.0$	1
ω	0.5	Equal weighting	1
ρ	0.03	Social discount rate	1
N	500	Firm count	2
p_0	0.7	Exploitation success rate	2
q_0	0.15	Exploration success rate	2
ν	0.3	Capability exponent	2
ξ	2.5	Pareto tail parameter	2
θ	0.5	Selection intensity	2

4.1 Comparative Statics

Sensitivity analysis reveals that the model's policy recommendations are most sensitive to:

1. **Window closure rate** γ: Faster-closing windows dramatically increase the shadow price of delay and shift α^* toward exploration.

2. **Lock-in speed** λ: Higher lock-in speed increases the urgency of present sovereignty investment, creating tension with the exploration imperative.

3. **Exploration returns** β: Higher returns to exploration make the case for reallocation stronger, lowering α^*.

5 Conclusion

The model formalises the essay's central argument: Europe's sovereignty challenge is not simply a question of securing today's technology stack, but of allocating scarce resources between exploitation and exploration under closing paradigm windows and self-reinforcing dependency. The three-layer structure reveals that:

1. The **optimal allocation** heavily favours exploration ($\alpha^* \approx 0.35$), suggesting current sovereignty spending is misallocated toward presentism.

2. **Designed constraint** ($\sigma \approx \sigma^*$) can raise exploration productivity, but only when absorptive capacity exceeds a threshold — programme design must take this seriously.

3. The **cost of delay** is non-linear and accelerating. Each year of inaction narrows the remaining paradigm windows, compounding the eventual cost of building alternatives.

Necessity is the mother of invention. But as this model shows, only *designed* necessity — calibrated to capability, timed to windows, and sustained by investment — can serve as the grandmother of sovereignty.