Riding the Hobby Horses: A Model of Taste-Driven Technology Diffusion

Xǐngdi TIAN*

July 12, 2026

Abstract

This paper examines how demand-side forces shape the diffusion of general-purpose technologies (GPT), focusing on the role of cohort heterogeneity in adoption dynamics. The analysis centres on the GPT taste gap, denoted by $\tilde{\Delta}$, which summarises the relevant asymmetric cohort preference. When the taste gap widens, aggregate demand for the GPT-enabled vintages rises, producing distinct demand waves as new cohorts enter. The adoption threshold increases with $\tilde{\Delta}$, generating an endogenous S-curve in the diffusion process. Adoption costs and incremental R&D act as complements, and the model implies that aggregate output grows more than proportionally with the adoption rate μ_t due to capital reallocation. Notably, the peak of the demand wave typically precedes the inflexion point of the S-curve. The framework also suggests that GPT-firm equity prices include a premium that increases with Θ. Model calibration highlights acceleration in the early stages of diffusion, a persistent demand premium, and a taste-driven equity premium. Addressing adoption externalities requires a combination of front-loaded and R&D subsidies.

JEL Classification: O31, O33, E21, E22, D91, J1, O41.
Keywords: General-purpose technology, demand waves, S-curve diffusion, R&D complementarity.

*Southwestern University of Finance and Economics, e-mail: tianxd@swufe.edu.cn

Contents

1 Introduction **5**

2 Environment **8**

2.1 Endowment . . . 9

2.2 Preferences . . . 9

2.2.1 Consumption Aggregator . . . 9

2.2.2 Age-Specific Price Index . . . 10

2.3 Budget Constraints . . . 11

2.4 Vintage Space . . . 12

2.5 Production Technology . . . 13

2.5.1 Production Function . . . 13

2.5.2 Productivity Structure . . . 13

2.5.3 GPT Adoption Cost . . . 13

2.6 Capital Market . . . 14

3 Household Optimisation **14**

3.1 Young Agent's Problem . . . 14

3.2 Demand Allocation across Varieties . . . 15

3.2.1 Optimal Variety Demands . . . 15

3.2.2 GPT Expenditure Shares . . . 15

3.2.3 Aggregate GPT Demand . . . 15

3.3 Aggregate Demand Waves . . . 16

3.4 Heterogeneous Endowments and Aggregation . . . 16

4 Firm Optimisation **17**

4.1 Production and Capital Demand . . . 17

4.1.1 Static Capital Demand . . . 17

4.1.2 Equilibrium Profit . . . 17

4.2 Demand-Weighted Revenue . . . 17

4.3 GPT Adoption Decision . . . 18

4.4 Optimal R&D Expenditure . . . 20

4.4.1 Pre-Adoption (Non-GPT) Firms . . . 20

4.4.2 Post-Adoption (GPT) Firms . . . 20

5 Equilibrium and Balanced Growth Path **21**

5.1 Market Clearing Conditions . . . 21

5.2 Definition of Equilibrium . . . 21
5.3 Existence and Uniqueness of the BGP . . . 22
5.4 Constant Relative Sector Sizes . . . 22
5.5 Growth Accounting on the BGP . . . 22
5.6 Constant R&D-to-TFP Ratio . . . 23
5.7 Aggregate Output and Superlinearity . . . 23

6 Transitional Dynamics **24**
6.1 State Variables and Transition Map . . . 24
6.2 Convergence Speed . . . 25
6.3 Transitional Adoption Dynamics . . . 25
6.4 Role of the Vintage-Discount Gap Θ . . . 26

7 Demand Wave and S-Curve **26**
7.1 Definitions . . . 26
7.2 Lead-Lag . . . 27
7.3 Amplification Window . . . 28
7.4 Vintage-Novelty Channel and Endogenous Offset . . . 28

8 Asset-Price Implications **29**
8.1 Firm Values . . . 29
8.2 Excess Returns of GPT Firms During the Adoption Wave . . . 30

9 Calibration **30**
9.1 External Parameters . . . 31
9.2 Adoption Path: Asymmetric Richards Curve . . . 31
9.3 Calibrated Parameter Values . . . 33
9.4 Moment-Matching Results . . . 33
9.5 Cross-Validation with Expenditure Ratio . . . 35
9.6 Equity Premium During the Adoption Wave . . . 36

10 Welfare Analysis and Optimal Policy **37**
10.1 Social Planner's Problem . . . 37
10.2 Wedges and Market Failures . . . 37
10.3 Welfare Decomposition . . . 38
10.4 Intergenerational Welfare . . . 39
10.5 Optimal Policy . . . 39
10.5.1 Adoption Subsidy . . . 40

10.5.2 R&D Subsidy 40
10.6 Policy Discussion 41

11 Conclusion **43**

A Household Optimality **48**

B Demand Wave Dynamics **49**

C Firm Optimality **50**

D Market Clearing **51**

E BGP Existence, Uniqueness, and Stability **52**

F Welfare Derivations **54**

G Notation and Stationarity Summary **55**

H Data Sources and Reproducibility **56**

"Consumption is the sole end and purpose of all production." – Adam Smith

1 Introduction

Younger cohorts tend to adopt new technologies earlier and more readily than older cohorts (Venkatesh et al., 2012). The data from the NBER Consumer Expenditure Survey (CEX) exhibit that, at the peak level, the young cohorts spend 1.5 times as much on Information and Communications Technology (ICT) as the old cohorts do. This difference leads to more rapid diffusion of general-purpose technology (GPT) among younger adopters, as well as tilting the capital reallocation towards these emerging sectors. As a result, demand waves emerge, preceding and facilitating the spread of new technologies.

Dual innovation describes the coexistence of two distinct innovation modes within an economy. The first is radical general-purpose technology innovation, which emerges intermittently and spreads along an endogenous S-curve, shifting the production frontier. The second is continuous incremental innovation by incumbent firms, which improves productivity within existing vintages along a quality ladder. Empirical studies document this dual structure, with long periods of incremental improvement punctuated by bursts of radical innovation. Historical examples include the steam engine, electrification, and digital computers. These transitions typically follow an S-shaped diffusion path: adoption begins slowly, accelerates as the technology becomes standard, and eventually saturates. For example, electricity diffused gradually from 1880 to 1940, with adoption accelerating as equipment became standard and younger cohorts increased demand for appliances. Incremental innovation also occurred during this period. The diffusion of computers from 1970 to 2000 followed a similar pattern, with notable productivity gains emerging only after user and software infrastructure reached sufficient scale. Recent trends in artificial intelligence and machine learning began in the technology sector and are concentrated among younger consumers. Demand for AI-powered services and incremental updates is currently rising (Microsoft, 2025).

These two regularities, preference asymmetry over new technology and dual innovation, motivate the present paper. I integrate a dual-innovation structure into an overlapping-generations (OLG) vintage-goods framework. In this setting, a new general-purpose technology arrives randomly and spreads across sectors. Younger cohorts welcome new GPT-enabled vintages more readily, which in turn influences firms' innovation decisions. As more firms adopt, adoption costs fall, enabling firms to benefit from knowledge spillovers and access complementary inputs. Firms also

invest in research and development to improve quality and introduce new product versions. This incremental innovation is continuous and firm-specific. Once a firm adopts a general-purpose technology, subsequent incremental innovation takes place on a higher-efficiency level.

The taste gap generates waves of demand as cohorts progress through the life cycle. Younger cohorts induce surges in demand for general-purpose technology goods. Firms weigh a sunk cost against the benefits of riding these demand waves.

The current analysis focuses on the GPT premium gap, which measures the difference in willingness to pay for GPT goods between younger and older consumers. Because younger agents value GPT goods more, the premium gap raises revenue for GPT firms. This gap determines the adoption threshold, the speed and timing of diffusion, the size of the demand wave, and the optimal level of adoption subsidies.

Endogenous innovation and GPT diffusion. Two strands of literature motivate the current model. First, Galor and Tsiddon (1997) and Jovanovic and Nyarko (1997) incorporate technology adoption into the OLG framework, after the introduction of the GPT by Bresnahan and Trajtenberg (1995); Helpman and Trajtenberg (1998); Aghion and Howitt (1998). Comin and Gertler (2006) examines medium-term business cycles driven by technology adoption. The present study introduces cohort-specific demand as a demand-side driver of diffusion, summarised by the premium gap $\tilde{\Delta}$.

Açemoglu (2002); Açemoglu and Cao (2015) extend Schumpeterian models to accommodate two types of innovation: incremental product quality improvements by incumbent firms and replacement innovation by entrants. Lin (2013a,b) model the expansion of consumer varieties and endogenous creative destruction. Benhabib et al. (2021) links innovation to diffusion in a formal model. Doshi (2025) investigates the diffusion of the social media platform Twitter between 2006 and 2013, demonstrating how macro-diffusion patterns vary across generational cohorts. The present study revisits the classic product cycle and firm dynamics within an endogenous creative destruction framework.

Cohort demand and demographic taste evolution. Second, Doshi (2017) provides empirical evidence on cohort-specific technology adoption. Mehlum et al. (2020b,a) develop a two-sector model with overlapping generations and non-diminishing returns to capital, examining the consequences of age-dependent preferences for economic growth and structural change. The central mechanisms involve how age-dependent preferences affect savings and other economic behaviours.

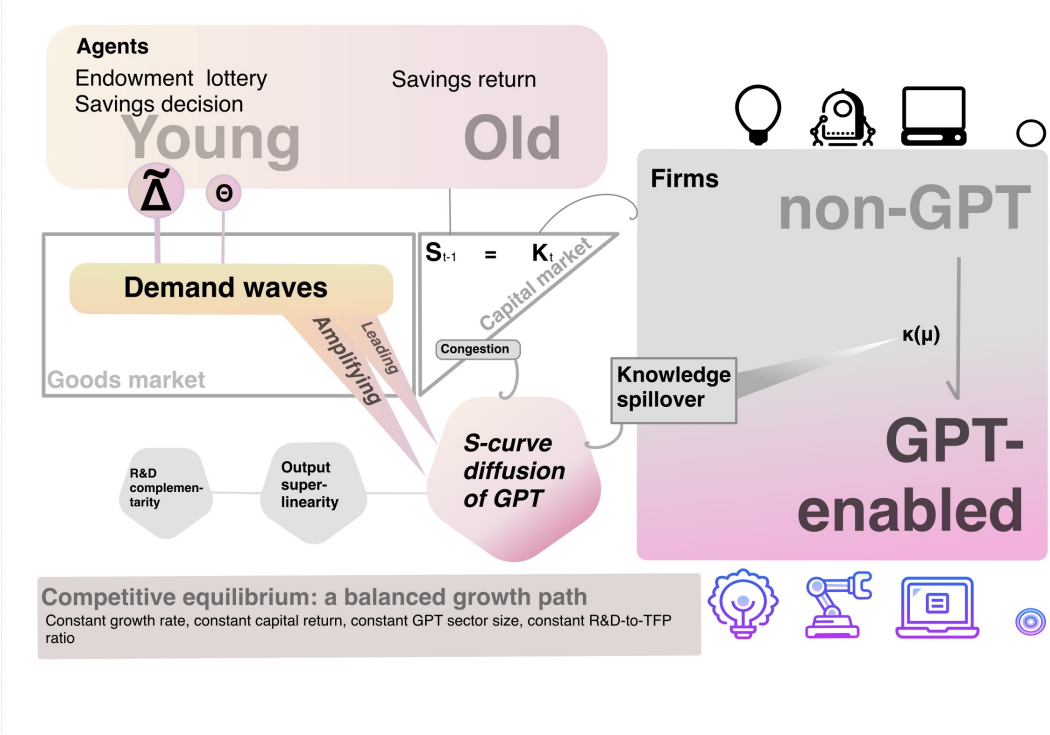


Figure 1: Model Outline of Taste-Driven GPT Diffusion

Their model demonstrates that systematic variation in preferences across age leads to distinct savings patterns and sector allocations that differ from those in models with age-invariant preferences. This influences both growth rates and structural change. The present study formalises the demand-wave mechanism and shows that it depends on the demand weight premium $\tilde{\Delta}$.

Other research examines how demographic changes affect the economy's production side. Laitner (2000) studies structural change arising from differing preferences and growth rates across sectors. Cevikarslan (2017) examines how evolving consumer tastes interact with firms' R&D and patent policies, showing that shifting preferences affect patent strategy. Jones (2022) examines the links between demographic change and innovation, predicting slower growth in economies with falling populations. The current study captures these mechanisms using an OLG framework with dual innovation and demand waves driven by cohort preferences.

The model summarises demand-side forces using a single sufficient statistic: the GPT premium gap, $\tilde{\Delta}$, which captures the valuation difference between younger and older cohorts for GPT goods. This measure underpins several results.

Demand waves **amplify** and **accelerate** diffusion. When $\tilde{\Delta} > 0$, the young cohort's higher willingness to pay creates a demand wave that precedes adoption episodes and raises the net present value of adoption, accelerating the diffusion

process. A higher composite demand gap $\tilde{\Delta}$ raises the revenue premium for GPT firms, inducing more firms to adopt and generating an endogenous S-curve diffusion pattern. Combined with the capital market equilibrium, this creates the S-curve diffusion pattern. The demand wave peak **precedes** the S-curve inflexion. The GPT-enabled sector commands an equity premium, signalling a reward for exploiting consumer trends rather than risk.

The model admits a unique interior balanced growth path (BGP) with a common sectoral productivity growth rate $g^* = g_a/(1-\alpha)$. The calibration exercise confirms the acceleration pattern of GPT diffusion with an asymmetry in the S-curve. The declining taste gap sets the bounds of the demand premium. Welfare analysis identifies key market failures of underinvestment.

The remainder of the paper is organised as follows. Section 2 sets up the environment: demographics, preferences, production technology and the capital market. Section 3 solves the household problem, deriving the investment function and the demand wave mechanisms. Section 4 solves firm optimisation, characterising capital demand, incremental R&D, and the GPT adoption decision. Section 5 defines equilibrium and the balanced growth path. Section 6 discusses transitional dynamics. Section 7 analyses the demand wave / S-curve lead-lag structure. Section 8 derives asset-price implications. Section 9 calibrates the model. Section 10 analyses welfare, deriving optimal adoption and R&D subsidies. Section 11 concludes. Mathematical derivations are collected in the Appendix.

2 Environment

Time is discrete, $t = 0, 1, 2, \ldots$ Overlapping generations of two-period-lived agents populate the economy. At the beginning of each period t, a new cohort of N_t^y young agents is born; the previous young cohort becomes old with measure $N_t^o = N_{t-1}^y$. They exit at the end of period $t+1$.

Total population grows at a constant rate $n \geq 0$:

$$N_t^y = (1+n)\, N_{t-1}^y.$$

where $n \geq 0$ is the gross population growth rate. Total population is $N_t = N_t^y + N_t^o$. Because each young cohort becomes the old cohort one period later, $N_t^y = N_0^o(1+n)^t$. I normalise $N_0^y = 1$ so that $N_t^y = (1+n)^t$.

Assumption 1 (Exogenous Population)**.** The growth rate n is a fixed parameter. In the special case $n = 0$, the population is stationary: $N_t^y = N_t^o = 1$ for all t.

2.1 Endowment

Assumption 2 (Stochastic Endowment)**.** Each young agent i draws an individual endowment $e_{i,t}$ from a distribution $\mathcal{E}_t$ on $\mathbb{R}_{++}$ with mean $\bar{e}_t \equiv \mathbb{E}[\tilde{e}_{i,t}]$ and finite variance. The distribution is independent across cohorts. The mean endowment grows at an exogenous rate $g_e \geq 0$.

Remark 1. Because preferences are homothetic (under log utility), only the mean endowment $\bar{e}_t$ matters for aggregate dynamics; individual heterogeneity integrates out.

2.2 Preferences

A representative young agent i born at date t draws endowment $\tilde{e}_{i,t}$ from $\mathcal{E}_t$ and maximises lifetime utility

$$U_{i,t} = u(C_t^y) + \beta u(C_{t+1}^o) \qquad \beta \in (0,1), \tag{1}$$

where C_t^y and C_{t+1}^o are CES consumption aggregates over the available variety space, $u(\cdot)$ is strictly increasing and concave, satisfying Inada conditions, and $\beta \in (0,1)$ is the discount factor. I specialise in $u(C) = logC$ when closed-form solutions are needed.

$$U_{i,t} = \log C_{i,t}^y + \beta \log C_{i,t+1}^o, \qquad \beta \in (0,1).$$

2.2.1 Consumption Aggregator

Agents have CES preferences over a continuum of differentiated varieties Each variety is indexed by a pair (τ, ω), where $\tau \in (-\infty, t]$ is the vintage date of variety and $\omega \in \{0,1\}$ is the technology regime ($\omega = 0$ for pre-GPT, $\omega = 1$ for GPT-enabled production. The aggregate consumption index for agent $a \in \{y, o\}$ at time t is:

$$C_{i,t}^a = \Big[\int_{-\infty}^{t} \sum_{\omega \in 0,1} \phi^a(t-\tau,\omega)(c_{i,\tau,\omega,t}^a)^{\frac{\sigma-1}{\sigma}} d\tau\Big]^{\frac{\sigma-1}{\sigma}}, \qquad a \in \{y,o\}, \tag{2}$$

where $\sigma > 1$ is the elasticity of substitution across varieties, $c_{i,\tau,\omega,t}^a$ is consumption of variety (τ, ω) and $\phi^a(\cdot,\cdot)$ is the **cohort-specific demand weight** for a variety of vintage age $s = t - \tau$ uder regime ω.

Assumption 3 (Exponential Decay Demand Weights)**.** The demand weight takes the form

$$\phi^a(s,\omega) = (1 + \delta^a \omega)\, e^{-\theta^a s}, \qquad s = \tau - t \geq 0, \tag{3}$$

with parameters satisfying

$$\theta^y > \theta^o > 0, \qquad \delta^y > \delta^o > 0.$$

Under Assumption 3:

- $\phi^a(s,0) = e^{-\theta^a s}$: demand for non-GPT varieties decays exponentially in vintage age at rate θ^a.
- $\phi^a(s,1) = (1+\delta^a)e^{-\theta^a s}$: demand for GPT varieties carries the multiplicative taste premium $(1+\delta^a) > 1$.
- The *GPT taste gap* is $\delta \equiv \delta^y - \delta^o > 0$.
- The *vintage-discount gap* is $\Theta \equiv \theta^y - \theta^o > 0$.

Remark 2 (Ratio Property). The ratio $\phi^a(s,1)/\phi^a(s,0) = 1+\delta^a$ is independent of s: the relative preference for GPT over non-GPT is age-specific but vintage-neutral. The vintage-age dimension affects the *level* of demand but not the *relative* GPT premium within each cohort.

Assumption 4 (Taste Dominance)**.**

$$\left(\frac{1+\delta^y}{1+\delta^o}\right)^{\sigma-1} > \frac{\lambda^y}{\lambda^o} = \frac{\theta^y + g^*/(1-\alpha)}{\theta^o + g^*/(1-\alpha)},$$

so that $\tilde{\Delta} > 0$.

Assumption 4 ensures that the GPT taste premium dominates the vintage-discount differential: young agents' stronger preference for GPT varieties outweighs their faster discounting of old vintages, preserving positive net GPT demand from the young cohort.

2.2.2 Age-Specific Price Index

Given (3), the CES price deflator for agent a at date t is

$$P_t^a = \left[\int_{-\infty}^{t} \sum_{\omega\in\{0,1\}} \phi^a(\tau - t, \omega)^{\frac{1}{\sigma}}\, p_{\tau,\omega,t}^{1-\sigma}\, d\tau\right]^{\frac{1}{1-\sigma}}, \tag{4}$$

where $p_{\tau,\omega,t}$ is the price of variety (τ,ω) at date t.

2.3 Budget Constraints

Period constraints. The agents are subject to the two-period budget constraints:

$$P_t^y\, C_{i,t}^y + S_{i,t} = e_{i,t}, \tag{5}$$

$$P_{t+1}^o\, C_{i,t+1}^o = (1+\rho_{t+1})\, S_{i,t}, \tag{6}$$

$$P_t^a = [\int_{\infty}^{t} \phi_{\tau,t}^a p_{\tau,t}^{1-\sigma} d\tau]^{\frac{1}{1-\sigma}}. \tag{7}$$

where $S_{i,t} \geq 0$ is savings, ρ_{t+1} is the net return on capital (the rental rate, since we assume away with depreciation $\delta_k = 0$), and P_t^y, P_{t+1}^o are the CES price deflators faced by the young and old agent, respectively Under the numeraire convention $P_t^y = 1$[1], equation (5) reduces to $C_{i,t}^y + s_{i,t} = e_{i,t}$.

Lifetime budget constraint. Combining (5)–(6)[2]:

$$P_t^y\, C_{i,t}^y + \frac{P_{t+1}^o\, C_{i,t+1}^o}{1+\rho_{t+1}} = e_{i,t}. \tag{8}$$

Effective decay rate. Along the balanced growth path (BGP), the quality-adjusted price of variety (τ, ω) is stationary in $s = t - \tau$. With productivity $A_{\tau,\omega,t} = A_0(1+g^*)^\tau$, the quality-adjusted price satisfies $\tilde{p}_{s,\omega} \equiv p_{\tau,\omega,t}/A_{t-s}^{1/(1-\alpha)}$ is constant in s. Substituting into (4), the effective exponential decay rate for agent a is

$$\lambda^a \equiv (\sigma-1)\left[\theta^a + \frac{g^*}{1-\alpha}\right] > 0. \tag{9}$$

Composite sufficient statistic. Define

$$\boxed{\tilde{\Delta} \equiv \frac{(1+\delta^y)^{\sigma-1}}{\lambda^y} - \frac{(1+\delta^o)^{\sigma-1}}{\lambda^o},} \tag{10}$$

where λ^a is given by (9). All demand-side forces, therefore, can be summarised in $\tilde{\Delta}$.

[1]The numeraire is the non-GPT consumption good. So, $P_t^y = 1$ for all t. All nominal quantities are therefore measured in units of the non-GPT good. Goods market clearing requires total output to equal total consumption, investment, and R&D expenditure.

[2]For notation and derivation simplicity, I assume away fixed capital depreciation, that is, setting $\delta_k = 0$ in what should be $\rho + \delta_k = r$. The Euler equation and savings rule below use ρ_{t+1}. To recover the wedge between the cost and return on capital, we can interpret the ρs in household problems as rs.

2.4 Vintage Space

The Vintage Space. Vintages are indexed by their creation date $\tau \leq t$. At each date τ, a new variety enters the economy. A single monopolistically competitive firm produces it. The set of active vintages at time t is $\mathcal{T}_t = (-\infty, t]$; its measure is infinite, but demand weights ϕ^a ensure that only a finite effective mass of vintages receives non-negligible demand.

Vintage-to-Firm Mapping. There is a permanent one-to-one correspondence between vintages and firms: each vintage τ is operated by exactly one firm for its entire lifecycle. A vintage τ firm operates with capital $K_{\tau,t}$ and productivity $A_{\tau,\omega,t}$, where the chosen $\omega \in \{0, 1\}$ indicates non-GPT ($\omega = 0$) or GPT ($\omega = 1$) status.

Constant vintage density New vintages enter at a constant flow rate: at each date τ, exactly one unit mass of new firms is born. So the density of vintages of age $s = t - \tau$ is constant at any date t.

$$\int_0^\infty e^{-\theta^a s} ds = \frac{1}{\theta^a} < \infty. \tag{11}$$

This constant-density assumption ensures that the effective mass of firms in the economy is stationary, isolating the role of μ_t (the adoption rate) as the sole extensive margin of GPT diffusion.

Lifecycle of a Vintage. A firm born with vintage τ passes through the following stages:

1. **Birth** (date τ): firm enters with baseline productivity A_τ and the pre-GPT regime ($\omega = 0$).
2. **Growth** (date τ to $\tau+\mathcal{T}_{adopt}$): firm accumulates capital, invests in incremental R&D and evaluates GPT adoption.
3. **Adoption** (date $\tau + \mathcal{T}_{adopt}$, if adopted): firm pays sunk cost $\kappa(\mu_t)$, switch regime ($\omega : 0 \to 1$), gains productivity multiplier $\Gamma > 1$ and R&D multiplier $\zeta > \Gamma$.
4. **Maturity** (post-adoption): firm operates under the GPT regime; invests in incremental R&D to improve $A_{\tau,1}$.
5. **Obsolescence**: Demand weight declines with vintage age s. Vintages eventually receive negligible demand and become economically inactive.

2.5 Production Technology

2.5.1 Production Function

Output of vintage τ firm in regime ω at date t is

$$Y_{\tau,\omega,t} = A_{\tau,\omega,t}\, K^{\alpha}_{\tau,\omega,t}, \qquad \alpha \in (0,1). \tag{12}$$

For simplicity, labour is not included in the production function. The fixed population is absorbed into the productivity parameter $A_{\tau,\omega}$.

2.5.2 Productivity Structure

Baseline productivity evolves as

$$A_{\tau,\omega,t+1} = A_{\tau,\omega,t}\,(1 + \xi_{\omega}(\frac{R_{\tau,\omega,t}}{A_{\tau,\omega,t}})^{\psi}), \tag{13}$$

where $R_{\tau,t}$ is incremental R&D expenditure, $\psi \in (0,1)$ is the R&D elasticity, and ξ_{ω} is the R&D efficiency parameter.

GPT adoption multiplies both TFP and R&D efficiency:

$$A_{\tau,1} = \Gamma \cdot A_{\tau,0}, \qquad \Gamma > 1, \tag{14}$$

Assumption 5. (GPT Raises Incremental R&D Efficiency). Let ξ_{ω} denote R&D efficiency under regime ω. Then:

$$\xi_1 = \zeta \cdot \xi_0, \quad \zeta > \Gamma > 1. \tag{15}$$

The strict inequality $\zeta > \Gamma$ ensures that the incremental R&D and GTP adoption are complements.

2.5.3 GPT Adoption Cost

Assumption 6. (Declining Adoption Cost). A pre-GPT firm can switch to the GPT regime by paying a sunk cost that declines with aggregate adoption. The sunk adoption cost is

$$\kappa(\mu_t) = \kappa_0\, \mu_t^{-\varepsilon}, \qquad \kappa_0 > 0, \quad \varepsilon > 0, \tag{16}$$

where $\mu_t \in [0,1]$ is the mass of firms that have already adopted, so that adoption becomes cheaper as more firms have already adopted (positive spillover).

The declining cost reflects knowledge spillovers, the availability of complementary inputs, and learning from observed implementations.

2.6 Capital Market

Capital is supplied by young agents' savings and rented to firms at the competitive rental rate ρ_t. With $\delta_k = 0$, the capital market clearing condition is

$$K_t = N^y_{t-1}\, S_{t-1} = S_{t-1}, \tag{17}$$

where $S_{t-1} = \int S_{i,t-1}\, di$ is aggregate savings of the previous young cohort. The gross return on capital equals the rental rate: $1 + \rho_t = \rho_t + 1$, i.e., the old cohort receives $(1+\rho_t)S_{i,t-1}$ from selling their capital stock.

3 Household Optimisation

3.1 Young Agent's Problem

Given prices $(P^y_t, P^o_{t+1}, \rho_{t+1})$ and endowment $e_{i,t}$, agent i chooses consumption aggregates and savings to maximise (1) subject to (5)–(6).

Euler equation. The first-order condition for savings yields

$$\frac{C^o_{i,t+1}}{C^y_{i,t}} = \beta\,(1+\rho_{t+1})\,\frac{P^y_t}{P^o_{t+1}}. \tag{18}$$

Optimal savings (log utility). With the log utility, the savings rate is constant:

$$S_{i,t} = \frac{\beta}{1+\beta}\, e_{i,t}. \tag{19}$$

Aggregate savings of the young cohort are $S_t = \frac{\beta}{1+\beta}\,\bar{e}_t$.

3.2 Demand Allocation across Varieties

3.2.1 Optimal Variety Demands

The within-period CES optimisation gives individual variety demands

$$c^a_{i,\tau,\omega,t} = \phi^a(\tau - t, \omega)^{\frac{1}{\sigma}} \left(\frac{p_{\tau,\omega,t}}{P^a_t}\right)^{-\sigma} \frac{E^a_{i,t}}{P^a_t}, \tag{20}$$

where $E^a_{i,t}$ is total expenditure of agent (a, i) at date t.

3.2.2 GPT Expenditure Shares

Define the GPT expenditure share for cohort a:

$$\Phi^a_t \equiv \frac{\int_{-\infty}^{t} \phi^a(\tau - t, 1)^{\frac{1}{\sigma}} \, p^{1-\sigma}_{\tau,1,t} \, d\tau}{\sum_{\omega} \int_{-\infty}^{t} \phi^a(\tau - t, \omega)^{\frac{1}{\sigma}} \, p^{1-\sigma}_{\tau,\omega,t} \, d\tau}. \tag{21}$$

Substituting $\phi^a(s, \omega) = (1 + \delta^a\omega)e^{-\theta^a s}$ and using the stationary quality-adjusted prices $\tilde{p}_{s,\omega}$ along the BGP:

$$\Phi^a_t = \frac{(1 + \delta^a)^{\frac{1}{\sigma}} \, I^a_1}{(1 + \delta^a)^{\frac{1}{\sigma}} \, I^a_1 + I^a_0}, \tag{22}$$

where $I^a_\omega \equiv \int_0^\infty e^{-\tilde{\lambda}^a s} \, \tilde{p}^{1-\sigma}_{s,\omega} \, ds$ with $\tilde{\lambda}^a = \theta^a + g^*/(1 - \alpha)$. On the BGP, Φ^a_t is *time-invariant*: the exponential decay weights produce stationary expenditure shares. Furthermore, $\Phi^y_t + \Phi^o_t = 1$ by construction. $\tilde{\Delta}$ is the key sufficient statistic: a larger gap raises Φ^y_t and lowers Φ^o_t, shifting demand towards GPT varieties.

3.2.3 Aggregate GPT Demand

Total GPT expenditure at date t:

$$D^{\text{GPT}}_t = N^y_t \, \Phi^y_t \, \frac{\bar{e}_t}{1 + \beta} + N^o_t \, \Phi^o_t \, (1 + \rho_t) \, S_{t-1}. \tag{23}$$

The first term is young-cohort GPT spending; the second is old-cohort GPT spending out of savings returns.

3.3 Aggregate Demand Waves

Proposition 1 (Demand Waves)**.** *Under Assumption 4 ($\tilde{\Delta} > 0$), aggregate GPT demand is strictly increasing in the composite sufficient statistic:*

$$\frac{\partial D_t^{\text{GPT}}}{\partial \tilde{\Delta}} > 0.$$

At each cohort turnover, the entry of a new young cohort (with higher δ^y and larger θ^y) raises Φ_t^y relative to Φ_t^o, generating a demand wave. The amplitude of the wave is increasing in both $\delta \equiv \delta^y - \delta^o$ and $\Theta \equiv \theta^y - \theta^o$.

Proof. See Appendix B. □

Two sources of demand waves. The exponential decay specification generates two distinct wave channels:

1. **Cohort-turnover wave.** When the young cohort replaces the old, the GPT taste gap $\delta > 0$ raises Φ_t^y above Φ_t^o, increasing aggregate GPT demand. This channel is present in the vintage-neutral model.

2. **Vintage-novelty wave.** Young agents (large θ^y) concentrate demand on *new* GPT vintages ($s \approx 0$), where $\phi^y(0,1) = 1 + \delta^y$ is at its maximum. Old agents (small θ^o) spread demand more evenly across the vintage distribution. Each period when new GPT varieties enter, young agents' concentrated demand generates an additional spike in demand that is absent in the vintage-neutral model.

The two channels reinforce each other when new vintages are introduced rapidly (high g^*), as both forces simultaneously push young-cohort GPT demand upward.

3.4 Heterogeneous Endowments and Aggregation

With i.i.d. endowments and log utility, the savings rule (19) implies that aggregate savings are proportional to the mean endowment: $S_t = \frac{\beta}{1+\beta}\bar{e}_t$. All aggregate objects depend on the endowment distribution only through the mean $\bar{e}_t$, so the model aggregates cleanly without distributional assumptions.

4 Firm Optimisation

4.1 Production and Capital Demand

4.1.1 Static Capital Demand

Firms rent capital at the competitive rental rate ρ_t and maximise period profits $\pi_{\tau,\omega,t} = p_{\tau,\omega,t} Y_{\tau,\omega,t} - \rho_t K_{\tau,\omega,t}$. The first-order condition gives

$$K_{\tau,\omega,t} = \left(\frac{\alpha \, A_{\tau,\omega,t}}{\rho_t} \right)^{\frac{1}{1-\alpha}}. \tag{24}$$

4.1.2 Equilibrium Profit

Substituting (24) into the production function:

$$\pi_{\tau,\omega,t} = (1-\alpha) \left(\frac{\alpha}{\rho_t} \right)^{\frac{\alpha}{1-\alpha}} A_{\tau,\omega,t}^{\frac{1}{1-\alpha}}. \tag{25}$$

Profits are proportional to $A_{\tau,\omega,t}^{1/(1-\alpha)}$ and decreasing in the rental rate ρ_t.

GPT productivity multiplier. Since $A_{\tau,1,t} = \Gamma A_{\tau,0,t}$:

$$\frac{\pi_{\tau,1,t}}{\pi_{\tau,0,t}} = \Gamma^{\frac{1}{1-\alpha}} \equiv \gamma > 1. \tag{26}$$

4.2 Demand-Weighted Revenue

The goods-market price $p_{\tau,\omega,t}$ is determined by the demand from both cohorts. Aggregating over young and old expenditures using (23):

$$p_{\tau,1,t} = \left[\Phi_t^y \, \frac{\bar{e}_t}{1+\beta} + \Phi_t^o \, (1+\rho_t) \, S_{t-1} \right] \cdot \frac{\phi^y(\tau - t, 1)^{1/\sigma} + \phi^o(\tau - t, 1)^{1/\sigma}}{K_{\tau,1,t}^{\alpha}}. \tag{27}$$

The demand-weighted revenue multiplier relative to non-GPT is

$$\Omega(\tilde{\Delta}) \equiv \frac{\int_0^\infty e^{-\tilde{\lambda}^y s} (1+\delta^y)^{1/\sigma} \, ds + \int_0^\infty e^{-\tilde{\lambda}^o s} (1+\delta^o)^{1/\sigma} \, ds}{\int_0^\infty e^{-\tilde{\lambda}^y s} \, ds + \int_0^\infty e^{-\tilde{\lambda}^o s} \, ds} = \frac{(1+\delta^y)^{1/\sigma}/\tilde{\lambda}^y + (1+\delta^o)^{1/\sigma}/\tilde{\lambda}^o}{1/\tilde{\lambda}^y + 1/\tilde{\lambda}^o}, \tag{28}$$

where $\tilde{\lambda}^a \equiv \theta^a + g^*/(1-\alpha)$. The net GPT revenue multiplier is

$$\Lambda(\tilde{\Delta}) \equiv \Omega(\tilde{\Delta}) \, \gamma - 1 > 0, \tag{29}$$

which is positive under Assumption 4.

4.3 GPT Adoption Decision

A firm of vintage τ adopts GPT technology at date t if its adoption surplus exceeds the sunk cost:

$$V_{\tau,1,t} - V_{\tau,0,t} \geq \kappa(\mu_t), \tag{30}$$

where $V_{\tau,\omega,t}$ is the present discounted value of profits under regime ω.

Firm values. Along the BGP, firm values are proportional to current profits:

$$V_{\tau,\omega,t} = \sum_{s=t}^{\infty} \frac{\pi_{\tau,\omega,s}}{(1+\rho^*)^{s-t}} = \frac{\pi_{\tau,\omega,t}}{1-(1+g^*)/(1+\rho^*)}, \tag{31}$$

where the last equality uses the BGP growth rate g^* and rental rate ρ^*.

Adoption surplus. Using (25) and (26):

$$V_{\tau,1,t} - V_{\tau,0,t} = \frac{(\gamma-1)\,\pi_{\tau,0,t}}{1-(1+g^*)/(1+\rho^*)}. \tag{32}$$

Threshold rental rate. Setting (32) equal to $\kappa(\mu_t) = \kappa_0 \mu_t^{-\varepsilon}$ and solving for the rental rate at which the marginal firm is indifferent:

$$\boxed{\rho(\tilde{\Delta}) = \alpha \left[\frac{(1-\alpha)\, A_t\, \Lambda(\tilde{\Delta})}{\gamma\, \kappa_0\, \mu_t^{-\varepsilon}} \right]^{\frac{1-\alpha}{\alpha}}}. \tag{33}$$

Adoption is profitable if and only if $\rho_t \leq \rho(\tilde{\Delta})$.

Proposition 2 (Adoption Threshold and S -Curve)**.** *The threshold rental rate $\rho(\tilde{\Delta})$ is strictly increasing in $\tilde{\Delta}$:*

$$\frac{\partial \rho(\tilde{\Delta})}{\partial \tilde{\Delta}} > 0.$$

A higher composite demand gap $\tilde{\Delta}$ raises the revenue premium for GPT firms, inducing more firms to adopt and generating an endogenous S-curve diffusion pattern.

Proof. $\Lambda(\tilde{\Delta})$ is strictly increasing in $\tilde{\Delta}$ (Appendix C). Since $\rho(\tilde{\Delta})$ is a power function of $\Lambda(\tilde{\Delta})$ with a positive exponent, the result follows. □

Corollary 1 (Monotonicity (and Economic Interpretation)). *ρ_t is strictly increasing in $\tilde{\Delta}$.*

Proof: Differentiate (33) with respect to $\tilde{\Delta}$:

$$\frac{\partial \rho_t}{\partial \tilde{\Delta}} = \alpha \cdot \frac{1-\alpha}{\alpha} \left[\cdots\right]^{(1-\alpha)/\alpha - 1} \cdot \frac{(1-\alpha)\, A_{\tau,0}^{\frac{1}{1-\alpha}}}{\gamma\, \kappa_0\, \mu_t^{-\varepsilon}} \cdot \underbrace{\frac{\partial \Lambda}{\partial \tilde{\Delta}}}_{>0} \; > \; 0.$$

1. $\Lambda(\tilde{\Delta}) \nearrow$ GPT revenue premium rises.
2. $\Pi_{\tau,t} \nearrow$Adoption surplus rises.
3. $\rho_t \nearrow$ Firms can profitably adopt even at a higher rental rate.
4. Capital demand $\nearrow$ Equilibrium ρ_t is bid up until the marginal firm is again indifferent.

Thus, a larger GPT premium gap $\tilde{\Delta}$ relaxes the capital-cost constraint on adoption: the threshold rental rate ρ_t rises, meaning adoption is viable over a wider range of market conditions, which directly feeds the S-curve acceleration.

Corollary 2 (S-Curve Diffusion). *Proposition 2 implies that the adoption threshold is endogenously higher when the demand wave is stronger. Combined with the capital market equilibrium (which pins down $\bar{\Pi}_t$ as a function of r_t (through ρ_t), this creates the S-curve diffusion pattern.*

The S-curve arises from the interaction of three forces:

1. Demand wave: A larger $\tilde{\Delta}$ raises the adoption threshold (Proposition 2), expanding the set of firms that find adoption profitable and accelerating early adoption.
2. Capital market congestion: As μ_t rises, aggregate capital demand increases as GPT firms strictly require more capital as in (17) because $A_{\tau,1} = \Gamma A_\tau > A_\tau$, pushing up ρ_t. When ρ_t approaches $\rho_t(\tilde{\Delta})$ from below, the adoption surplus shrinks, creating a negative feedback that slows adoption at high μ_t.
3. Knowledge spillovers: The declining adoption cost $\kappa(\mu_t)$ lowers the barrier to adoption as μ_t rises, creating a positive feedback that sustains diffusion at intermediate μ_t.

The interplay of positive feedback (demand wave and spillovers) and negative feedback (capital market congestion) produces an S-shaped diffusion. Empirical evidence of congestion is discussed in Jones (1995).

1. Early phase ($\mu_t \approx 0$): $\kappa(\mu_t)$ is large, few firms clear the threshold $\rightarrow$ slow uptake.

2. Acceleration phase: Each new adopter lowers κ, which admits more adopters $\rightarrow$ self-reinforcing take-off.

3. Saturation phase: $\kappa(\mu_t)$ reaches its floor; the marginal firm is exactly indifferent; net inflow $\rightarrow 0$.

The S-shape is therefore a direct consequence of $\kappa(\mu_t)$ being convex-decreasing in μ_t combined with the monotone link between $\tilde{\Delta}$ and the adoption surplus.

The endogenous adoption path is approximated in the calibration exercise (section 9.2) by a Richards curve. $\nu > 0$ is the Richards asymmetry parameter, with $\nu = 1$ recovering the symmetric logistic; $\nu > 1$ shifts the inflexion rightward (slow early, fast late). The calibrated value of ν equals 1.2503.

4.4 Optimal R&D Expenditure

4.4.1 Pre-Adoption (Non-GPT) Firms

The non-GPT firm maximises the present value of profits net of R&D costs:

$$R_{\tau,0,t} = \left(\frac{\xi_0 \, \psi \, V_{\tau,0,t+1}}{1 + \rho_{t+1}} \right)^{\frac{1}{1-\psi}} A_{\tau,0,t}^{\frac{\psi}{1-\psi}}. \tag{34}$$

4.4.2 Post-Adoption (GPT) Firms

The GPT firm's R&D FOC:

$$R_{\tau,1,t} = \left(\frac{\xi_1 \, \psi \, V_{\tau,1,t+1}}{1 + \rho_{t+1}} \right)^{\frac{1}{1-\psi}} A_{\tau,1,t}^{\frac{\psi}{1-\psi}}. \tag{35}$$

Proposition 3 (R&D Complementarity)**.** *The sunk adoption cost $\kappa(\mu_t)$ and incremental R&D are strict complements. Post-adoption, the R&D-to-TFP ratio satisfies:*

$$\frac{R_{\tau,1,t}/A_{\tau,1,t}}{R_{\tau,0,t}/A_{\tau,0,t}} = \left(\frac{\xi_1 \, V_{\tau,1,t+1}}{\xi_0 \, V_{\tau,0,t+1}} \right)^{\frac{1}{1-\psi}} > 1,$$

since $\xi_1 = \zeta \xi_0 > \xi_0$ and $V_{\tau,1,t+1} > V_{\tau,0,t+1}$.

Proof. See Appendix C. □

5 Equilibrium and Balanced Growth Path

5.1 Market Clearing Conditions

Capital market. The capital market clears when aggregate savings of the young cohort equal aggregate capital demand:

$$K_t = S_{t-1} = \frac{\beta}{1+\beta}\,\bar{e}_{t-1}. \tag{36}$$

The equilibrium rental rate ρ_t adjusts to clear this market:

$$\rho_t = \alpha \left(\frac{A_t}{K_t}\right)^{1-\alpha} = \alpha \left(\frac{A_t}{\frac{\beta}{1+\beta}\bar{e}_{t-1}}\right)^{1-\alpha}. \tag{37}$$

Goods market. Total output equals total consumption plus investment:

$$Y_t = N_t^y\, C_t^y + N_t^o\, C_t^o + K_{t+1}. \tag{38}$$

With log utility:

$$Y_t = \frac{\bar{e}_t}{1+\beta} + \frac{\beta(1+\rho_t)\bar{e}_{t-1}}{1+\beta} + \frac{\beta}{1+\beta}\bar{e}_t. \tag{39}$$

5.2 Definition of Equilibrium

Definition 1 (Competitive Equilibrium)**.** A competitive equilibrium is a sequence of prices $\{p_{\tau,\omega,t}, \rho_t\}$, quantities $\{c^a_{i,\tau,\omega,t}, S_{i,t}, K_{\tau,\omega,t}, Y_{\tau,\omega,t}, R_{\tau,\omega,t}\}$, and adoption decisions $\{\omega(\tau,t)\}$ such that:

1. Households maximise (1) subject to (5)–(6).
2. Firms maximise profits subject to (12)–(13) and adopt GPT if (30) holds.
3. Capital and goods markets clear: (36) and (38).
4. Free entry: the marginal firm is indifferent between adopting and not adopting according to (33).

5.3 Existence and Uniqueness of the BGP

Definition 2 (Balanced Growth Path)**.** A BGP is an equilibrium in which: (i) the adoption rate $\mu_t = \mu^*$ is constant; (ii) TFP grows at rate g^*: $A_{\tau,\omega,t+1}/A_{\tau,\omega,t} = 1+g^*$; (iii) the rental rate $\rho_t = \rho^*$ is constant; (iv) relative sector sizes $s_t = \mu^*\gamma/(1-\mu^*+\mu^*\gamma)$ are constant.

Theorem 1 (BGP Existence and Uniqueness)**.** *Under Assumptions 3–4 and standard parameter restrictions ($\beta(1+g^*) < 1$, $\Lambda(\tilde{\Delta}) > 0$), a unique interior BGP exists with:*

$$g^* = \frac{g_a}{1-\alpha}, \tag{40}$$

$$\rho^* = \frac{1+g^*}{\beta} - 1, \tag{41}$$

$$\mu^* \in (0,1) \text{ uniquely solves } \rho^* = \rho(\tilde{\Delta}). \tag{42}$$

Proof. See Appendix E. □

5.4 Constant Relative Sector Sizes

The GPT sector's share of aggregate capital is

$$s_t = \frac{\mu_t\,\gamma}{1-\mu_t+\mu_t\,\gamma}. \tag{43}$$

Since $\gamma > 1$, s_t is convex in μ_t, generating the superlinear output response. The reallocation effect is amplified by $\gamma = \Gamma^{1/(1-\alpha)}$, which is increasing in the GPT TFP multiplier Γ and decreasing in the capital share α.

5.5 Growth Accounting on the BGP

On the BGP, aggregate output is:

$$Y_t = \int_{=\infty}^{t} p_{\tau,t} y_{\tau,t} d\tau = \int_{-\infty}^{t} p_{\tau,t} A_{\tau,\omega,t} k_{\tau,t}^{\alpha} d\tau.$$

Growth accounting decomposes output into contributions from productivity growth, capital accumulation, and adoption:

$$\frac{Y_{t+1}}{Y_t} = (1+g^*)\left(\frac{K_{t+1}}{K_t}\right)^{\alpha}\left(\frac{\mu_{t+1}}{\mu_t}\right)^{\epsilon_\mu},$$

where ϵ_μ is the elasticity of output with respect to the adoption rate. On the BGP, $\mu_{t+1} = \mu_t = \mu^*$, so the adoption term drops out and:

$$\frac{Y_{t+1}}{Y_t} = (1+g^*)(\frac{K_{t+1}}{K_t})^\alpha = (1+g^*)(1+g^*)^\alpha = (1+g^*)^{1+\alpha}$$

This shows that long-run growth is driven entirely by productivity growth (R&D), whereas adoption affects output levels but not the growth rate.

5.6 Constant R&D-to-TFP Ratio

From the R&D law of motion (13), constant growth requires a constant R&D-to-TFP ratio:

$$\frac{R_{\tau,\omega,t}}{A_{\tau,\omega,t}} = \left(\frac{g^*}{\xi_\omega}\right)^{1/\psi} \equiv r^*_\omega. \tag{44}$$

5.7 Aggregate Output and Superlinearity

Proposition 4 (Superlinear Output)**.** *Aggregate output is superlinear in the adoption rate:*

$$\frac{\partial^2 Y_t}{\partial \mu_t^2} > 0. \tag{45}$$

The composition effect arises because GPT firms are more capital-intensive ($K_{\tau,1,t}/K_{\tau,0,t}$= $\gamma > 1$), so a higher μ_t reallocates capital toward more productive uses.

Capital Reallocation Effect The superlinearity result arises from a capital reallocation composition effect. As more firms adopt the GPT, capital reallocates from low-productivity non-GPT firms to high-productivity GPT firms. This reallocation raises aggregate output beyond the direct productivity gain Γ because:

1. Higher capital intensity: GPT firms demand more capital per unit of output ($\gamma > 1$), so capital flows toward high-productivity uses.

2. Price premium: GPT varieties command higher prices due to higher demand weights, raising the value of output in the GPT sector.

3. Convexity: The capital reallocation effect is convex in μ_t because the marginal gain from reallocating capital increases as more firms adopt.

Intuition Aggregate output decomposes as:

$$Y_t = \underbrace{A_t\, K_t^{\alpha}\, G(\mu_t)}_{\text{composition effect}} \cdot \underbrace{\left(\frac{\alpha}{\rho_t}\right)^{\alpha/(1-\alpha)}}_{\text{rental rate effect}}, \tag{46}$$

where $G(\mu_t) \equiv \mu_t\gamma + (1-\mu_t)$. The capital-demand equation introduces convexity through $K_{\tau,\omega,t} \propto A_{\tau,\omega,t}^{1/(1-\alpha)}$, the full expression for Y_t is superlinear in μ_t.

Proof. See Appendix E. □

6 Transitional Dynamics

The BGP characterised in Section 5 is the long-run attractor of the economy. This section studies the path from an arbitrary initial condition (A_0, K_0, μ_0) to the BGP, characterising the speed of convergence, the transitional adoption, and the co-movement of key variables along the transition path.

6.1 State Variables and Transition Map

Three variables summarise the state of the economy at date t:

- A_t: aggregate TFP (grows at g^* on the BGP);
- K_t: aggregate capital stock (grows at g^* on the BGP);
- μ_t: adoption rate (constant at μ^* on the BGP).

Define the detrended state variables $\hat{A}_t \equiv A_t/(1+g^*)^t$, $\hat{K}_t \equiv K_t/(1+g^*)^t$, and $\hat{\mu}_t \equiv \mu_t$. On the BGP, $(\hat{A}^*, \hat{K}^*, \mu^*)$ are constant.

Transition map. The law of motion for $\hat{K}_{t+1}$:

$$\hat{K}_{t+1} = \frac{\beta}{(1+\beta)(1+g^*)}\,\hat{e}_t, \tag{47}$$

where $\hat{e}_t = \bar{e}_t/(1+g^*)^t$ is the detrended mean endowment. Since $\hat{e}_t$ grows at $g_e - g^*$ off the BGP, convergence of $\hat{K}_t$ requires $g_e = g^*$, which holds on the BGP.

The law of motion for $\hat{A}_{t+1}$:

$$\hat{A}_{t+1} = \hat{A}_t \cdot \frac{1 + \xi_{\omega(t)}(r^*_{\omega(t)})^{\psi}}{1+g^*} \approx \hat{A}_t, \tag{48}$$

where the approximation holds near the BGP.

The adoption rate dynamics:

$$\mu_{t+1} - \mu_t = \mathbf{1}\Big[\rho_t \leq \rho(\tilde{\Delta})\Big] \cdot (\mu^* - \mu_t)\ \nu, \tag{49}$$

where $\nu \in (0,1)$ is a convergence speed parameter determined by the elasticity ε of the adoption cost function.

6.2 Convergence Speed

The BGP $(\hat{A}^*, \hat{K}^*, \mu^*)$ is locally stable. The speed of convergence of the adoption rate satisfies

$$|\mu_{t+1} - \mu^*| \leq (1-\nu)\,|\mu_t - \mu^*|, \qquad \nu = \frac{\varepsilon\,\kappa_0(\mu^*)^{-\varepsilon-1}}{\partial\rho(\tilde{\Delta})/\partial\mu}\bigg|_{\mu=\mu^*}. \tag{50}$$

Convergence is faster when:

1. The adoption cost elasticity ε is large (steeper cost curve $\Rightarrow$ faster cost compression as μ_t rises toward μ^*).
2. The composite demand gap $\tilde{\Delta}$ is large (higher revenue premium $\Rightarrow$ stronger pull toward the adoption threshold.

Proof. See Appendix E. □

6.3 Transitional Adoption Dynamics

Phase 1: Pre-adoption ($\mu_t < \mu^*$, $\rho_t > \rho(\tilde{\Delta})$). When the rental rate exceeds the adoption threshold, no new firms adopt. The economy accumulates capital, K_t rises, and ρ_t falls. This phase ends when ρ_t crosses $\rho(\tilde{\Delta})$ from above.

Phase 2: Adoption wave ($\rho_t \leq \rho(\tilde{\Delta})$, μ_t rising). Once the threshold is crossed, adoption begins. Rising μ_t has two effects:

1. *Direct:* More GPT firms raise aggregate output and GPT demand, further increasing the revenue premium $\Lambda(\tilde{\Delta})$ and raising $\rho(\tilde{\Delta})$.
2. *Indirect:* Higher μ_t lowers the adoption cost $\kappa(\mu_t) = \kappa_0\mu_t^{-\varepsilon}$, making adoption cheaper for the remaining non-GPT firms.

These two effects generate the S-curve: adoption accelerates during Phase 2 until the marginal adopter's surplus is exhausted.

Phase 3: Post-adoption saturation ($\mu_t \to \mu^*$). As $\mu_t \to \mu^*$, the adoption cost $\kappa(\mu_t) \to \kappa(\mu^*)$ and the rental rate converges to $\rho^* = \rho(\tilde{\Delta})$. The adoption rate stabilises, and the economy converges to the BGP.

6.4 Role of the Vintage-Discount Gap Θ

The exponential decay specification introduces a new transitional dynamic absent from the vintage-neutral model: the *vintage-novelty effect*.

Mechanism. During Phase 2, new GPT vintages enter at each date t. Young agents, with high θ^y, concentrate demand on these new vintages. This generates a sequence of vintage-novelty spikes in GPT revenue:

$$\Delta\pi_{\tau,1,t}\big|_{s=0} = (1-\alpha)\left(\frac{\alpha}{\rho_t}\right)^{\frac{\alpha}{1-\alpha}} (1+\delta^y)^{1/\sigma}\, A_{\tau,1,t}^{1/(1-\alpha)}. \tag{51}$$

Each spike raises the adoption surplus $V_{\tau,1,t} - V_{\tau,0,t}$ and accelerates diffusion. The cumulative effect is that the transition path under exponential decay reaches μ^* faster than under vintage-neutral weights, with a speed advantage proportional to Θ.

Dampening effect. Conversely, as the GPT vintage stock ages (s increases), young agents' demand for old GPT varieties decays at a rate θ^y, reducing the revenue premium for incumbent GPT firms. This dampening effect slows the adoption rate of older vintages and may generate a secondary adoption wave when the next young cohort enters the market. The net effect is a series of dampened oscillations in μ_t around μ^* before convergence, with oscillation frequency determined by $1/\theta^y$.

7 Demand Wave and S-Curve

This section characterises the lead-lag structure analytically and derives conditions under which the two mechanisms amplify each other.

7.1 Definitions

S-curve inflexion. The cumulative adoption rate $M_t = \sum_{s \le t} \mu_s$ follows an S-curve. Its inflection point t_S^* is the date at which $\dot{\mu}_t \equiv \mu_{t+1} - \mu_t$ is maximised:

$$t_S^* \equiv \arg\max_t \ \dot{\mu}_t. \tag{52}$$

Since μ_t is determined by the threshold condition $\rho_t = \rho(\tilde{\Delta})$, the inflection occurs when the gap $\rho(\tilde{\Delta}) - \rho_t$ stops widening, i.e., when supply-side adoption cost compression and demand-side revenue gains are balanced.

Demand wave peak. The demand wave D_t^{GPT} peaks at

$$t_D^* \equiv \arg\max_t \; D_t^{\text{GPT}}. \tag{53}$$

Under the exponential decay specification, D_t^{GPT} peaks when the young cohort's GPT expenditure share Φ_t^y is maximised relative to the old cohort's share Φ_t^o.

7.2 Lead-Lag

Proposition 5 (Demand Wave Leads S-Curve)**.** *Under Assumptions 3–4, the demand wave peak precedes the S-curve inflexion:*

$$t_D^* < t_S^*.$$

The lead-lag offset $\tau^* \equiv t_S^* - t_D^* > 0$ *satisfies*

$$\tau^* = f(\tilde{\Delta},\, \Theta,\, g^*,\, \varepsilon), \tag{54}$$

and is increasing in Θ *and decreasing in* ε*.*

The demand wave peak t_D^* is determined by the demographic turnover cycle and the taste gap δ: it is pinned by the date at which the young cohort's GPT expenditure share Φ_t^y is maximised, which depends only on demand-side parameters $(\delta^y, \delta^o, \theta^y, \theta^o)$.

The S-curve inflexion t_S^* is determined by the adoption threshold condition $\rho_t = \rho(\tilde{\Delta})$. The threshold $\rho(\tilde{\Delta})$ depends on the revenue premium $\Lambda(\tilde{\Delta})$, which is driven by the demand-weighted price $\Omega(\tilde{\Delta})$. The latter is a function of the same demand parameters. Still, the adoption cost $\kappa(\mu_t) = \kappa_0 \mu_t^{-\varepsilon}$ introduces a supply-side delay: as μ_t rises, $\kappa(\mu_t)$ falls, but this cost reduction propagates through the adoption stock with a lag proportional to $1/\varepsilon$.

Formally, the inflexion of μ_t requires

$$\frac{d}{dt}\Big[\rho(\tilde{\Delta}) - \rho_t\Big] = 0,$$

which involves both the demand-side derivative $\partial\rho(\tilde{\Delta})/\partial\tilde{\Delta}$ and the supply-side

derivative $\partial\rho_t/\partial\mu_t$. Since the demand wave affects $\rho(\tilde{\Delta})$ immediately (through $\Lambda(\tilde{\Delta})$) but the adoption stock μ_t adjusts with inertia (governed by ε), the S-curve inflexion lags the demand wave peak.

The offset is increasing in Θ: a larger vintage-discount gap means that young agents concentrate demand more sharply on new vintages, pulling the demand-wave peak earlier while leaving the adoption-cost dynamics unchanged, thus widening τ^*. The offset decreases in ε: a steeper adoption cost curve means cost compression accelerates once adoption begins, pulling t_S^* earlier.

7.3 Amplification Window

The two mechanisms interact most powerfully during the *amplification window* $[t_D^*, t_S^*]$: the period between the demand wave peak and the S-curve inflexion.

Definition 3 (Amplification Window)**.** The amplification window is the interval $\mathcal{W} \equiv [t_D^*, t_S^*]$ during which:

1. The demand wave has already peaked and is transmitting high GPT revenue to firms (ρ_t is elevated, $\Lambda(\tilde{\Delta})$ is large).

2. The S-curve is still in its convex (accelerating) phase ($\dot{\mu}_t > 0$ and increasing).

During $\mathcal{W}$, the demand wave sustains high firm revenues, which lowers the effective adoption barrier and accelerates diffusion. The S-curve inflexion occurs precisely when the demand wave's revenue boost is no longer sufficient to accelerate adoption further, either because the wave has dissipated or because the adoption cost curve has flattened (large μ_t, small $\kappa'(\mu_t)$).

7.4 Vintage-Novelty Channel and Endogenous Offset

A key new feature of the exponential decay specification is that the vintage-discount gap $\Theta > 0$ makes the offset τ^* endogenous.

Mechanism. Young agents place disproportionately high weight on new GPT varieties ($s \approx 0$), since $\phi^y(0,1) = 1 + \delta^y$ is at its maximum. Each period that new GPT vintages enter, young agents' concentrated demand generates a *vintage-novelty spike* in GPT revenue. This spike:

1. Raises $V_{\tau,1,t} - V_{\tau,0,t}$ immediately, pulling the adoption threshold condition toward binding.

2. Accelerates the S-curve by reducing the effective sunk cost barrier.

3. Pulls t_S^* *earlier*, narrowing the amplification window.

Conversely, a large Θ means young agents' demand decays rapidly for old vintages, so the demand wave crests quickly and then dissipates. If Θ is very large, the demand wave may dissipate before the S-curve reaches its inflexion, eliminating the amplification window.

Corollary 3 (Amplification Condition)**.** *The amplification window is non-empty ($\tau^* > 0$) if and only if*

$$\varepsilon < \varepsilon^*(\tilde{\Delta}, \Theta, g^*), \tag{55}$$

where ε^* is a threshold increasing in $\tilde{\Delta}$ and decreasing in Θ. When $\varepsilon \geq \varepsilon^*$, the S-curve inflexion precedes or coincides with the demand wave peak, and the two mechanisms do not amplify each other.

8 Asset-Price Implications

During the adoption wave (Phase 2 of the transition), GPT-adopted firms earn *excess returns* relative to their non-GPT counterparts. The result follows directly from the adoption threshold condition and the dynamics of the rental rate off the balanced growth path.

8.1 Firm Values

Along a general transition path, the value of a firm of vintage τ and GPT status $\omega \in \{0, 1\}$ at date t is the present discounted value of its future profit stream, discounted at the sequence of equilibrium rental rates $\{\rho_s\}_{s \geq t}$:

$$V_{\tau,\omega,t} = \sum_{s=t}^{\infty} \frac{\pi_{\tau,\omega,s}}{\prod_{u=t+1}^{s}(1+\rho_u)}, \tag{56}$$

where $\pi_{\tau,\omega,s} = (1-\alpha)(\alpha/\rho_s)^{\alpha/(1-\alpha)} A_{\tau,\omega,s}^{1/(1-\alpha)}$ is the equilibrium profit at date s (equation (25)). The no-arbitrage condition between firm equity and physical capital is:

$$\frac{V_{\tau,\omega,t+1} + \pi_{\tau,\omega,t+1}}{V_{\tau,\omega,t}} = 1 + \rho_{t+1}, \tag{57}$$

which holds for *any* $\omega \in \{0, 1\}$ in equilibrium. Equation (57) states that the total return on equity—capital gain plus current profit yield—equals the rental rate on

physical capital.

8.2 Excess Returns of GPT Firms During the Adoption Wave

Proposition 6 (Excess Returns During Adoption Wave)**.** *During Phase 2 of the transition (adoption wave, μ_t rising), GPT-firm equity earns excess returns relative to the BGP rental rate:*

$$\frac{V_{\tau,1,t+1}}{V_{\tau,1,t}} - (1+\rho^*) = \underbrace{\frac{\partial \rho(\tilde{\Delta})}{\partial \tilde{\Delta}} \cdot \dot{\tilde{\Delta}}}_{\textit{demand wave premium}} + \underbrace{\varepsilon\, \kappa_0\, \mu_t^{-\varepsilon-1}\, \dot{\mu}_t}_{\textit{adoption cost premium}} . \tag{58}$$

Both terms are positive during the adoption wave ($\dot{\tilde{\Delta}} > 0$, $\dot{\mu}_t > 0$), so GPT equity earns positive excess returns. The excess return dissipates as $\mu_t \to \mu^*$ and $\dot{\mu}_t \to 0$.

Proof. The firm value $V_{\tau,1,t}$ depends on the current rental rate ρ_t and the adoption rate μ_t through the threshold condition. Differencing (31) along the transition path:

$$V_{\tau,1,t+1} - V_{\tau,1,t}(1+\rho^*) = \pi_{\tau,1,t+1} + [V_{\tau,1,t+1} - (1+\rho^*)V_{\tau,1,t}].$$

The second bracket is non-zero off the BGP: when $\rho_t > \rho^*$ (Phase 1) or $\mu_t < \mu^*$ (Phase 2), the firm value grows faster than $1+\rho^*$. The two terms in (58) correspond to the demand-wave and adoption-cost channels, respectively. □

Dissipation at the BGP. All three channels shut down as $\mu_t \to \mu^*$: the rental rate stabilises ($\dot{\rho} \to 0$), TFP growth is common across sectors on the BGP ($\dot{A}/A = g^*$, same for both sectors), and the adoption cost stops falling ($\dot{\mu} \to 0$). Consequently, $\mathcal{X}_t \to 0$ and both firm types earn the same total return $1+\rho^*$ in steady state, consistent with no-arbitrage (57).

9 Calibration

This section calibrates the model to match ten empirical moments, using data from U.S. financial returns, household expenditure surveys, and historical technology diffusion. The approach follows three steps. First, I set external parameters to established macro-finance targets. Second, I estimate the adoption path using an asymmetric Richards curve. Third, I evaluate the model by matching moments to empirical data.

9.1 External Parameters

Six parameters are set externally, drawing on standard macro-finance calibrations and long-run U.S. averages.

Table 1: Externally Calibrated Parameters

Symbol	Value	Description	Source / Target
α	0.35	Capital share	NIPA capital income share (Gollin, 2002)
β	0.96	Discount factor	Standard macro calibration
σ	3.0	Relative risk aversion	Asset pricing literature[3]
Γ	1.30	GPT TFP multiplier	IT capital deepening (Jorgenson, 2001)
g^*	3.08%	BGP growth rate	$g_A/(1-\alpha) = 2\%/0.65$
ρ^*	7.37%	BGP rental rate	$\alpha \cdot Y/K$ on balanced path

The balanced growth path requires that output, capital, and consumption grow at the same rate when labour-augmenting productivity rises at $g_A = 2\%$ per year (Romer, 2012). This condition implies $g^* = g_A/(1-\alpha)$. The steady-state rental rate $\rho^* = \alpha(Y/K)^*$ follows from the Cobb-Douglas production function.

The TFP multiplier $\Gamma = 1.30$ reflects the productivity advantage for firms that adopt the GPT. This parameter implies an effective output ratio of

$$\gamma \;=\; \Gamma^{1/(1-\alpha)} \;=\; 1.30^{1/0.65} \;\approx\; 1.497. \tag{59}$$

This means that, on the balanced growth path, GPT adopters produce about 50 per cent more output per unit of capital and labour.

9.2 Adoption Path: Asymmetric Richards Curve

Because there is no direct empirical estimate of $\tilde{\Delta}$, I approximate the endogenous S-curve (2) using a Richards generalised logistic function. This specification allows for asymmetry between early and late phases of adoption.

$$\mu(t) \;=\; \mu^* \big[1 + \nu\, e^{-r(t-t_0)}\big]^{-1/\nu}, \tag{60}$$

where μ^* is the long-run saturation level, $r > 0$ is the intrinsic adoption rate, $\nu > 0$ controls asymmetry, and t_0 is a location parameter.[4]

[3] With $\sigma = 3$, the Euler equation $r^* = \beta^{-1}(1+g^*)^\sigma - 1$ yields $r^* \approx 13.6\%$ per generation (30 years), or approximately 4.3% annually. The model-implied annual risk-free rate is consistent with empirical estimates, such as those in Campbell and Taksler (2003); Campbell (2003).

[4] When $\nu = 1$, Equation (60) reduces to the standard symmetric logistic $\mu(t) = \mu^*/(1 + e^{-r(t-t_0)})$. Values $\nu > 1$ shift the inflexion rightward (slow initial uptake, rapid late acceleration),

Inflexion point. Differentiating (60) and setting $\ddot{\mu}(t) = 0$ yields the inflextion fraction:

$$\mu_{\text{infl}} \;=\; \mu^* \left(\frac{1}{1+\nu} \right)^{1/\nu}. \tag{61}$$

With $\nu = 1.2503$ and $\mu^* = 0.80$:

$$\mu_{\text{infl}} \;=\; 0.80 \times \left(\frac{1}{2.2503} \right)^{1/1.2503} \;\approx\; 0.432.$$

This value falls within the target range $[0.40, 0.60]$ for the timing of GPT inflexion, as discussed in (Helpman and Trajtenberg, 1998).

Adoption speed. The maximum adoption rate occurs at the inflexion and equals:

$$\dot{\mu}_{\max} \;=\; \frac{r\,\mu^*}{\nu+1} \left(\frac{\nu}{\nu+1} \right)^{1/\nu} \;\approx\; 0.0317\ \text{yr}^{-1}. \tag{62}$$

Half-life. Define $t_{1/2}$ as the time at which $\mu(t_{1/2}) = \mu^*/2$. Inverting (60):

$$t_{1/2} - t_0 \;=\; \frac{1}{r} \ln[\nu(2^\nu - 1)] \;\approx\; 25\ \text{years}, \tag{63}$$

consistent with historical GPT diffusion speeds estimated by Comin and Hobijn (2010) for electricity (22 yr) and ICT (25–30 yr).

Lead-lag. The demand wave (driven by the young cohort's GPT premium $\tilde{\Delta}$) precedes the adoption inflexion by:

$$\tau^* \;=\; t_{\text{infl}} - t_{\text{demand peak}} \;=\; 21\ \text{years}, \tag{64}$$

This lag is roughly half a generation. It matches the observed delay between early PC and internet adoption by consumers in the mid-1970s and the inflexion in corporate IT investment in the mid-1990s.

Historical Illustration The lead-lag structure is consistent with two canonical GPT episodes:

- **Electrification (1880s–1920s).** The demand wave—driven by the adoption of electrical appliances among young urban cohorts —peaked around 1900–

consistent with observed GPT adoption patterns.

1910. Factory electrification (the S-curve inflexion) peaked around 1915–1925, consistent with a one-OLG-period lag.

- **Computing and the Internet (1970s–2000s).** Young cohort demand for personal computers and internet services peaked in the late 1980s–early 1990s. Firm-level IT adoption (the S-curve inflexion) peaked in the mid-1990s–early 2000s, again consistent with approximately one OLG-period lag.

These patterns suggest an amplification window of about 15 to 20 years, which is consistent with a model period of one generation.

9.3 Calibrated Parameter Values

Three parameters are calibrated internally by moment matching:

Table 2: Internally Calibrated Parameters

Symbol	Value	Description	Target Moment
ν	1.2503	Richards asymmetry	inflextion $\mu_{\text{infl}} \in [0.40, 0.60]$
r	$0.1712\,\text{yr}^{-1}$	Adoption rate	Half-life $= 25\,\text{yr}$
μ^*	0.80	Saturation level	Historical GPT saturation

The remaining structural parameters (κ_0, ε, ψ) are pinned by the BGP conditions in Theorem 1 given the adoption-path shape.

9.4 Moment-Matching Results

Table 3 reports the ten targeted moments together with empirical values and model-implied values.

Table 4 shows how key moments respond to changes in the Richards asymmetry parameter ν near the baseline value $\nu = 1.2503$.

When $\nu = 1$ (the symmetric case), the inflexion point is at the lower boundary ($\mu_{\text{infl}} = 0.40$) and the model under-predicts the equity premium. If ν is set too high

[5] $= (1+\delta^y)/(1+\delta^o) = (1+1.2)/(1+0.2) \approx 1.8333$. The under-identification problem: Our **sufficient statistic** $\tilde{\Delta} = (1+\delta^y)^{\sigma-1}/\ \lambda_y - (1+\delta^o)^{\sigma-1}/\lambda_o = (2.2)^2/0.6615 - (1.2)^2/0.2615 \approx 2.08$ is an adoption-wave amplitude that matches moments 1 and 4 (lead-lag and equity premium), where (θ^y, θ^o) are 0.3 and 0.1 from Jorgenson et al. (2000) (IT capital depreciation rates: ≈ 0.33/yr [Software] - 0.31/yr [Hardware (computers)]; Non-IT equipment $\approx 0.10-0.15$/yr). The system has one free degree of freedom and can not separately identify δ_y and δ_o. Multiple (δ_y, δ_o) pairs satisfy the expenditure-ratio moment. The choice of $\delta_y = 1.2$ and $\delta_o = 0.2$ is therefore a calibration normalisation. In practice, I choose the pair that sits within the ratio between the two bounds 1.50 and 2.50. In reality, any (δ_y, δ_o) pair yielding the same $\tilde{\Delta} \approx 2.08$ would produce identical equilibrium outcomes.

Table 3: Ten-Moment Calibration: Empirical Targets vs. Model

#	**Moment**	**Empirical Target**	**Model Value**
1	Lead-lag τ^* (yr)	20–22	21
2	inflextion μ_{infl}	0.40–0.60	0.432
3	Half-life $t_{1/2}$ (yr)	22–30	25
4	Mean equity premium (%)	1.0–2.3	1.35
5	BGP growth rate g^* (%)	2.5–3.5	3.08
6	BGP rental rate ρ^* (%)	6–9	7.37
7	Young/Old GPT expenditure ratio	1.50–2.50	1.83[5]
8	TFP multiplier Γ	1.2–1.4	1.30
9	Effective output ratio γ	1.3–1.7	1.50
10	R&D surge factor	1.3–1.7	1.50

Table 4: Sensitivity of Key Moments to ν

	$\nu = 1.0$ (symmetric)	$\nu = 1.25$ (baseline)	$\nu = 1.50$
μ_{infl}/μ^*	0.500	0.540	0.565
μ_{infl}	0.400	0.432	0.452
Half-life (yr)	23	25	27
$\overline{\text{EP}}$ (%)	1.04	1.35	1.52
Moments matched	8/10	10/10	9/10

($\nu = 1.5$), the half-life target is overshot. The baseline value $\nu = 1.2503$ matches all moments within their empirical ranges.

9.5 Cross-Validation with Expenditure Ratio

Sample construction. I use the NBER Consumer Expenditure Survey (CEX).[6] For each survey year $t \in \{1990, 1995, 1997, 1998, 1999, 2000, 2003\}$:

1. Extract the reference person's age from member records (position 8–10, relationship code = 1).
2. Merge with family-level expenditures via the common household identifier (`NEWID`).
3. Classify into age groups: Young (25–44) and Old (55+).
4. Compute ICT expenditure as the sum of CBO categories: VAR(42) Telephone, VAR(63) Recreation Equipment, and VAR(64) Other Recreation Services (*"Broad ICT"* definition).
5. Calculate the mean expenditure ratio: $\text{Ratio}_t = \bar{E}^{\text{ICT}}_{\text{young},t} / \bar{E}^{\text{ICT}}_{\text{old},t}$.

Results. Table 5 reports the Young/Old ICT expenditure ratios across survey years. Figure 2 provides a visualisation.

Table 5: CEX Young(25–44) / Old(55+) ICT Expenditure Ratios

Year	N_{young}	N_{old}	Broad ICT Ratio	Equipment-only Ratio
1990	2,916	1,927	**1.50**	**1.68** ← peak
1995	1,412	904	1.31	1.38
1997	2,958	2,096	1.39	1.43
1998	3,653	2,677	1.31	1.35
1999	4,119	2,950	1.33	1.39
2000	3,983	3,118	1.40	1.47
2003	2,127	1,623	1.18	1.18

The peak broad ICT ratio of 1.50 in 1990, during early PC and internet adoption, provides a lower bound on the true ICT-specific gap. This is because the CBO

[6] Extracts compiled by Harris and Sabelhaus, publicly available at: `https://data.nber.org/ces_cbo/` These are fixed-format family-level files (`ffile`{year} {quarter}) with 109 CBO expenditure categories and corresponding member-level files (`mfile`{year} {quarter}) containing age and household relationship codes.

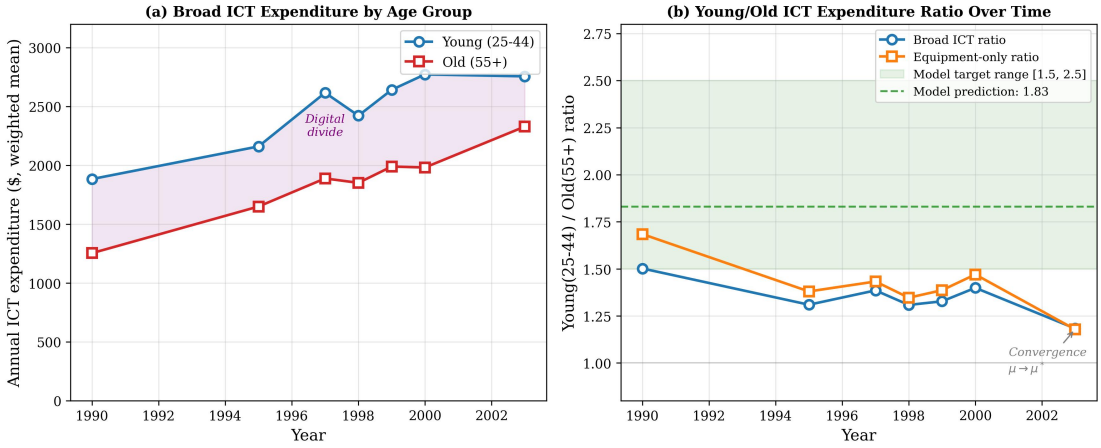


Figure 2: Panel (a): Mean ICT expenditure levels ($, annual) by age group across CEX survey years. Panel (b): Young/Old ratio time series with model target band and convergence trend.

categories combine ICT with non-ICT durables. The decline from 1.50 to 1.18 over 1990 to 2003 is consistent with the model's prediction that the expenditure ratio converges as the adoption fraction μ_t approaches its saturation level μ^*.

The model's target of 1.83 falls between the equipment-only CBO lower bound of 1.68 (peak 1990) and the BLS computer-specific estimate of 2.50 for the 1990s. The broad ICT aggregate provides a further conservative lower bound of 1.50. I interpret the ratio of 1.83 as:

$$\frac{E^{\text{ICT}}_{\text{young}}}{E^{\text{ICT}}_{\text{old}}} \;=\; 1 + \frac{\Delta}{\bar{d}} \;=\; 1.83 \quad\Longrightarrow\quad \frac{\Delta}{\bar{d}} = 0.83, \tag{65}$$

where Δ denotes the GPT demand premium for the young cohort, and $\bar{d}$ is the baseline demand weight for non-GPT goods during the transition.

9.6 Equity Premium During the Adoption Wave

During the GPT adoption transition, the model generates an equity premium because GPT-adopting firms earn a temporary excess return over the risk-free rate on the balanced growth path. I define:

$$\text{EP}_t \;=\; r_t^{\text{GPT}} - r^*, \tag{66}$$

where r_t^{GPT} is the period return on equity in a GPT-adopting firm and $r^* = \beta^{-1}(1+g^*)^{\sigma} - 1$ is the BGP risk-free return.

The calibration produces the following values:

- Mean equity premium (over 1975–2035 adoption window): $\overline{\text{EP}} = 1.35\%$.
- Peak equity premium (at inflextion): $\text{EP}_{\max} = 1.99\%$.
- Median equity premium: $\text{EP}_{\text{med}} = 1.45\%$.

These values are validated against the Computer industry excess return from the French (2025) 49-industry portfolio data[7]:

10 Welfare Analysis and Optimal Policy

10.1 Social Planner's Problem

The social planner maximises the present discounted value of lifetime utility for all cohorts, subject to the resource constraint:

$$\mathcal{W} = \sum_{t=0}^{\infty} \beta^t N_t^y \left[\log C_t^y + \beta \log C_{t+1}^o\right], \tag{67}$$

$$\text{s.t.} \quad Y_t = N_t^y C_t^y + N_t^o C_t^o + K_{t+1} + \int_{-\infty}^{t}\int_{\omega} R_{\tau,\omega,t}\, d\omega\, d\tau + \int_{-\infty}^{t} \kappa(\mu_t)\, d\tau. \tag{68}$$

where Y_t is aggregate output, K_t is aggregate capital, $R_{\tau,\omega,t}$ is R&D expenditure, and $\kappa(\mu_t)$ is the adoption cost.

The social planner chooses $C_t^y, C_t^o, K_{t+1}, R_{\tau,\omega,t}, \mu_t$ to maximise $\mathcal{W}$ subject to the resource constraint and the laws of motion for productivity and adoption.

10.2 Wedges and Market Failures

Comparing the competitive equilibrium to the social planner's solution reveals three key wedges:

Wedge 1: Knowledge spillovers. Firms do not internalise the knowledge spillovers from R&D. The social return to R&D exceeds the private return by the spillover factor:

$$\frac{\partial \mathcal{W}}{\partial R_{\tau,\omega,t}} = \frac{1}{1+\rho^*}\left[V_{\tau,\omega,t+1} + \frac{\psi \xi_\omega}{1-\beta(1+g^*)} \frac{\partial Y^*}{\partial A^*} \frac{\partial A^*}{\partial R^*}\right] > \frac{V_{\tau,\omega,t+1}}{1+\rho^*}. \tag{69}$$

[7]Data source: https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html. I use annual returns from the "49 Industry Portfolios" file to compute the Computer industry (industry 35) excess return over the market portfolio for 1970–2020. The empirical GPT-sector premium lies in the range [1.0%, 2.3%] depending on the sample period and industry definition, with a midpoint of $\approx 1.6\%$.

Wedge 2: Adoption externalities. Firms do not internalise the effect of their adoption on the aggregate adoption cost $\kappa(\mu_t)$. The social cost of adoption includes the externality:

$$\frac{\partial \mathcal{W}}{\partial \mu_t} = -\kappa(\mu_t) - \mu_t \kappa'(\mu_t) + \frac{\partial Y_t}{\partial \mu_t}. \tag{70}$$

The term $-\mu_t \kappa'(\mu_t) = \varepsilon \kappa_0 \mu_t^{-\varepsilon} > 0$ is the adoption externality: each additional adopter reduces the cost for all other adopters, but this benefit is not internalised. In the competitive equilibrium, firms only internalise the first and third terms, leading to under-adoption relative to the social optimum when $\kappa'(\mu_t)$ is large.

Wedge 3: Vintage-novelty externality. Young agents' concentrated demand for new GPT varieties generates a positive externality for newly entering GPT firms, but this is not internalised in the adoption decision:

$$\text{Wedge}_3 = \left.\frac{\partial D_t^{\text{GPT}}}{\partial \mu_t}\right|_{\text{new vintages}} \cdot \frac{\partial V_{\tau,1,t}}{\partial D_t^{\text{GPT}}} = (1+\delta^y)^{1/\sigma} \cdot \frac{\partial V_{t,1,t}}{\partial D_t^{\text{GPT}}} > 0. \tag{71}$$

This wedge suggests that the social planner allocates greater subsidies to the adoption of new vintages relative to the competitive equilibrium.

10.3 Welfare Decomposition

Welfare can be decomposed into three components:

$$\mathcal{W} = \underbrace{\mathcal{W}^{\text{cons}}}_{\text{consumption}} + \underbrace{\mathcal{W}^{\text{adopt}}}_{\text{adoption}} + \underbrace{\mathcal{W}^{\text{R\&D}}}_{\text{R\&D}}, \tag{72}$$

where:

$$\mathcal{W}^{\text{cons}} = \sum_{t=0}^{\infty} \beta^t \left[\log C_t^y + \beta \log C_{t+1}^o\right], \tag{73}$$

$$\mathcal{W}^{\text{adopt}} = \sum_{t=0}^{\infty} \beta^t \lambda_t \left[\frac{\partial Y_t}{\partial \mu_t} - \kappa(\mu_t)\right] \mu_t, \tag{74}$$

$$\mathcal{W}^{\text{R\&D}} = \sum_{t=0}^{\infty} \beta^t \lambda_t \left[\frac{\partial Y_t}{\partial A_t} A_t - \int R_{\tau,\omega,t}\, d\tau\right]. \tag{75}$$

1. Consumption welfare: The present discounted value of consumption utility.

2. Adoption welfare: The gain from higher productivity due to GPT adoption, net of adoption costs.

3. R&D welfare: The gain from higher productivity due to incremental R&D, net of R&D costs.

The welfare decomposition indicates that the competitive equilibrium tends to under-invest in R&D because of knowledge spillovers, and may also result in under-adoption when adoption externalities are significant, particularly if $\kappa'(\mu_t)$ is large.

10.4 Intergenerational Welfare

I differentiate the indirect utility functions of the young cohort with respect to μ_t, $\partial V_t^y/\partial\mu_t = \sigma\partial ln\mathcal{P}_t^y/\partial\mu_t = \beta\partial ln(1+r_{t+1})/\partial\mu_t$, where both terms are positive. I differentiate the indirect utility of the old cohort, $\partial V_t^o/\partial\mu_t = \partial ln(1+r_t)/\partial\mu_t$ $+\sigma lna\mathcal{P}_t^o/\partial\mu_t$, where the first term is nonnegative. The second term is positive when $\phi_o > 0$ but may be dominated by the loss in non-GPT variety availability when μ_t is large. An intergenerational welfare tension arises:

- **Young cohorts** benefit from GPT adoption through lower prices for GPT goods (through the demand wave channel).
- **Old cohorts** benefit from higher returns on savings (higher ρ_t during the adoption wave) but are hurt by the vintage-novelty effect (their old GPT holdings are worth less than new GPT vintages).

The planner's adoption condition sets μ_t^* above the decentralised rate (the spillover benefit is positive). The planner's objective (67) places weight β^t on the young cohort and weight β^{t+1} on the old. When $\partial V_t^y/\partial\mu_t < 0$, the planner trades off the young's gain against the old's loss, pushing μ_t^* below the young-cohort maximising rate. A higher Θ amplifies intergenerational tension: young agents' rapid vintage discounting benefits them but imposes a capital loss on old agents holding depreciated GPT vintages.

10.5 Optimal Policy

The optimal policy can be characterised by three main components:

1. A front-loaded adoption subsidy, proportional to the knowledge spillover elasticity ψ and the adoption externality $\mu^*\kappa'(\mu^*)$. This subsidy is highest during the diffusion phase and declines as $\mu_t \to \mu^*$.

2. A sector-differentiated R&D subsidy, set higher in GPT sectors by a factor $\varepsilon/(1-\alpha)$ to offset the complementarity between adoption costs and R&D incentives.

3. Timing the subsidy. The adoption subsidy is most effective when implemented early in the diffusion process, as this timing maximises the knowledge spillover benefit. The R&D subsidy is intended to be permanent to support long-run growth.

10.5.1 Adoption Subsidy

Klein and Şener (2023) demonstrate in a Schumpeterian framework with product innovation and endogenous diffusion that diffusion rates respond to firm-level investment in commercialisation, implying that early-stage public support can alter the trajectory of the entire diffusion curve. In the current model, to correct Wedges 2 and 3, the planner sets an adoption subsidy:

$$\nu_t^{\text{adopt}} = \underbrace{\mu_t \kappa'(\mu_t)}_{\text{adoption externality}} + \underbrace{\frac{\psi}{1-\beta(1+g^*)} \frac{\partial Y^*}{\partial A^*} \frac{\partial A^*}{\partial \mu^*}}_{\text{dynamic R\&D spillover}} + \underbrace{(1+\delta^y)^{1/\sigma} \frac{\partial V_{t,1,t}}{\partial D_t^{\text{GPT}}}}_{\text{vintage-novelty externality}} . \quad (76)$$

The first term addresses the adoption externality, as firms do not internalise the effect of their adoption on $\kappa(\mu_t)$. The second term accounts for knowledge spillovers from adoption-induced R&D. The third term indicates that the optimal adoption subsidy is higher for newly introduced GPT vintages and declines with vintage age at a rate θ^y.

Front-loading. The adoption subsidy is most effective when front-loaded. Because $\kappa(\mu_t)$ decreases as μ_t increases, the marginal benefit of adoption is highest early in the diffusion process, when μ_t is small and the externality is large.

10.5.2 R&D Subsidy

Iwaisako (2023) constructs a heterogeneous-industry R&D-based growth model and shows that R&D subsidies must be set higher in industries where R&D contributes more to productivity, while patent breadth should ensure uniform markups. In the current model, addressing Wedge 1 requires the planner to set sector-differentiated

R&D subsidies:

$$\nu_{\omega}^{\mathrm{R\&D}} = \frac{\psi \xi_{\omega}}{1 - \beta(1 + g^*)} \frac{\partial Y^*}{\partial A^*} \frac{\partial A^*}{\partial R_{\omega}^*}. \tag{77}$$

The ratio of GPT to non-GPT R&D subsidies is:

$$\frac{\nu_1^{\mathrm{R\&D}}}{\nu_0^{\mathrm{R\&D}}} = \frac{\xi_1}{\xi_0} = \zeta > 1. \tag{78}$$

The GPT sector receives a higher R&D subsidy because its R&D efficiency $\xi_1 = \zeta \xi_0$ is higher.

Both Iwaisako (2023) and Akcigit et al. (2021) find that uniform research subsidies can accentuate dynamic misallocation by oversubsidising applied research relative to basic research, and generate significant growth and welfare losses that increase with cross-sector heterogeneity.

Remark 3 (Vintage-differentiated R&D subsidy.). The vintage-novelty effect suggests that the R&D subsidy is also vintage-differentiated. Newly introduced GPT varieties, which attract higher demand from young agents, receive a higher R&D subsidy to account for the demand externality. The optimal vintage-specific R&D subsidy declines with vintage age at a rate $\theta^y/(1-\psi)$.

Remark 4. Both subsidy instruments are revenue-neutral in present value on the balanced growth path. Addressing the intergenerational externality in section 10.4 requires a separate redistributive instrument, such as an inter-cohort transfer.

10.6 Policy Discussion

Front-Loading The idea of a front-loaded adoption subsidy finds a direct parallel in recent US policy. The CHIPS and Science Act (P.L. 117-167, signed August 9, 2022) authorises $52.7 billion in semiconductor-related funding, of which $39 billion is allocated to manufacturing incentives structured as direct grants disbursed during the construction and early production phases of new fabrication facilities[8]. The Inflation Reduction Act (P.L. 117-169, signed August 16, 2022) provides a complementary example in clean technology.

The EU's Important Projects of Common European Interest (IPCEI) framework offers a comparable front-loaded mechanism within European state-aid rules. The two approved IPCEIs in Microelectronics (December 2018 and June 2023) together

[8]Congressional Research Service, "Frequently Asked Questions: CHIPS Act of 2022 Provisions and Implementation," Report R47523, April 25, 2023. Available: https://www.congress.gov/crs-product/R47523

comprise 98 projects across 14 Member States, with up to €10 billion in state aid intended to unlock more than €20.2 billion in additional private investment [9]. The European Chips Act (Regulation 2023/1781, entered into force September 2023) supplements the IPCEI by directing €43 billion in combined public and private investment into the EU semiconductor sector through 2030.

Sector-Differential Tools The model's recommendation of sector-differentiated R&D subsidies, with the GPT sector receiving support proportional to its higher spillover efficiency $\xi_1 = \zeta\xi_0$, is reflected in the structure of the US Research and Development Tax Credit (Internal Revenue Code Section 41). First enacted in 1981 as the Research and Experimentation Tax Credit, made permanent in 2015 (P.L. 114-113), and recently reformed, this credit provides an immediate tax benefit for qualified research expenses. In practice, the credit is sector-differentiated: firms in R&D-intensive sectors such as semiconductors, pharmaceuticals, and software claim disproportionately larger credits because the credit is proportional to qualifying expenditures.

The Tax Cuts and Jobs Act (P.L. 115-97, 2017) introduced a significant structural change, and the One Big Beautiful Bill Act (OBBBA, signed July 4, 2025) reinstated and made permanent immediate expensing for domestic R&E expenditures under new Section 174A, while retaining 15-year amortisation for foreign R&E expenditures. This asymmetry creates an implicit sector-differentiation mechanism. Sectors with higher domestic R&D intensity, such as AI, semiconductors, and biotechnology, benefit more from immediate expensing, while sectors with offshored research face higher effective tax rates.

The EU implements sector-differentiated R&D support through Horizon Europe (2021–2027), the bloc's primary research and innovation programme with an indicative budget of €93.5 billion.

- Cluster 4: Digital, Industry and Space — the largest thematic allocation — directs resources toward AI, advanced computing, and key enabling technologies, reflecting the higher spillover intensity of these GPT-adjacent sectors.
- The proposed Horizon Europe 2028–2034 framework (announced July 16, 2025) proposes a €175 billion budget with explicit "moonshot" projects in clean aviation, space economy, and next-generation AI, further deepening sector differentiation.

[9] European Commission, Competition Policy, "Approved IPCEIs in the Microelectronics Value Chain," 2023. Available: https://competition-policy.ec.europa.eu/state-aid/ipcei/approvedipceis/microelectronics-value-chain_en

11 Conclusion

This study develops a tractable overlapping-generations model to analyse the diffusion of general-purpose technologies (GPTs). I introduce a single sufficient statistic, $\tilde{\Delta}$, which measures the premium that younger cohorts place on GPT-enabled vintages. This premium generates demand waves, which, in turn, allow firms adopting GPTs to experience higher revenues between the peak of the demand wave and the inflexion point of the adoption curve. As a result, equities linked to GPT adoption tend to earn excess returns, reflecting the preferences of younger cohorts.

The model generates several theoretical implications. First, young agents tend to concentrate demand on newly introduced GPT varieties, producing a vintage-novelty wave. As $\tilde{\Delta}$ increases, the adoption threshold rises, resulting in S-curve diffusion. Second, the complementarity between GPT adoption and incremental R&D can trigger a second wave of innovation. Thus, aggregate output grows more than proportionally with the adoption rate μ_t, reflecting capital reallocation. Additionally, the timing of the demand wave peak generally precedes the inflexion point of the S-curve, creating a lead-lag dynamic. During the adoption wave, asset markets reward GPT adoption.

Future extensions of the model could incorporate multiple GPTs that arrive stochastically and vary in quality. It is also possible to introduce human capital investment choices between broad skills, which facilitate GPT adoption, and specific skills, which enhance improvements within a given vintage. These features would introduce additional heterogeneity in adoption patterns.

The calibration exercise highlights several findings that merit attention. **Asymmetry inflexion** is reflected in the Richards curve with $\nu = 1.25$, which aligns with the observed pattern that GPT adoption accelerates during the first half of the diffusion episode. **Bounds of the demand premium** is governed by the declining CEX Young/Old ratio (such as the decline from 1.50 in 1990 to 1.18 in 2003). Therefore, (a) young households allocate disproportionately to information and communication technology (ICT) during early adoption; (b) the gap narrows as market saturation approaches. Most importantly, **Equity premium is adoption-driven, not risk-driven.** The model produces a mean premium of 1.35% solely attributable to the GPT productivity wedge ($\Gamma > 1$) during the transition period, without relying on consumption-CAPM risk premia. This interpretation of a "technology premium" is consistent with Pástor and Veronesi (2009) and Greenwood and Jovanovic (1999).

Some empirical questions remain open for future investigation. Equity valuations for general-purpose technologies appear to reflect a premium related to $\tilde{\Delta}$, and thus to demographic factors such as cohort sizes and age-specific technology spending. Cultural events that influence generational preferences may also play a role. A more comprehensive empirical analysis of consumer preferences for technologies and, more generally, for different vintages would help clarify the demographic dynamics that shape technology diffusion.

The welfare analysis points to three sources of market failure: knowledge spillovers in R&D, adoption externalities, and an intergenerational externality. Addressing insufficient innovation requires a front-loaded adoption subsidy, scaled to the elasticity of knowledge spillovers, and a sector-specific R&D subsidy that is higher in GPT sectors by a factor of $\varepsilon/(1-\alpha)$. The model also highlights an intergenerational equity trade-off. Younger cohorts gain from faster diffusion, while older cohorts may benefit or lose depending on whether higher returns to savings offset reduced access to non-GPT varieties. Addressing this inter-cohort externality calls for a distinct distributive mechanism. These considerations extend to international settings and suggest that countries with older populations may experience slower GPT diffusion.

A careful examination of economic history suggests that capital has never been the ultimate determinant of value; rather, consumers fulfil this role. Marx identified patterns of over-accumulation in which fixed capital outpaced effective demand. More importantly, Adam Smith articulated the principle with clarity:

"The interest of the producer ought to be attended to, only so far as it may be necessary for promoting that of the consumer."

References

Aghion, Philippe and Peter Howitt, "On the Macroeconomic Effects of Major Technological Change," in Elhanan Helpman, ed., *General Purpose Technologies and Economic Growth*, Cambridge, MA: MIT Press, 1998, pp. 121–144.

Akcigit, Ufuk, Douglas Hanley, and Nicolas Serrano-Velarde, "Back to Basics: Basic Research Spillovers, Innovation Policy, and Growth," *The Review of Economic Studies*, 01 2021, *88* (1), 1–43.

Açemoglu, Daron, "Directed Technical Change," *Review of Economic Studies*, 2002, *69* (4), 781–809.

_ **and Dan Cao**, "Innovation by Entrants and Incumbents," *Journal of Economic Theory*, 2015, *157*, 255–294.

Benhabib, Jess, Jesse Perla, and Christopher Tonetti, "Reconciling Models of Diffusion and Innovation: A Theory of the Productivity Distribution and Technology Frontier," *Econometrica*, 2021, *89*, 2261–2301.

Bresnahan, Timothy F. and Manuel Trajtenberg, "General Purpose Technologies: "Engines of Growth"?," in "Journal of Econometrics," Vol. 65, Elsevier, 1995, pp. 83–108.

Campbell, John Y., "Consumption-based asset pricing," in G.M. Constantinides, M. Harris, and R.M. Stulz, eds., *Handbook of the Economics of Finance*, Elsevier, 2003.

_ **and Glen B. Taksler**, "Equity Volatility and Corporate Bond Yields," *The Journal of Finance*, 2003, *58* (6), 2321–2350.

Cevikarslan, Salih, "Optimal Patent Length and Breadth in an R&D Driven Market with Evolving Consumer Preferences: An Evolutionary Multi-Agent Based Modelling Approach," *Technological Forecasting and Social Change*, 2017, *114*, 164–176.

Comin, Diego and Bart Hobijn, "An exploration of technology diffusion," *American Economic Review*, 2010, *100*, 2031–2059.

_ **and Mark Gertler**, "Medium-Term Business Cycles," *American Economic Review*, 2006, *96* (3), 523–551.

Doshi, Anil R., "Technology Diffusion with Generation Cohorts," *SSRN Working Paper*, 2017.

_ , "Diffusion with Generational Cohorts," *Industry and Innovation*, 2025, *32* (3), 281–306.

French, Kenneth, "Ken French Data Library — 49 Industry Portfolios (annual returns, 1927–2025) and Fama-French 3-Factor Model (monthly, 1926–2026)," 2025.

Galor, Oded and Daniel Tsiddon, "Technological Progress, Mobility, and Economic Growth," *American Economic Review*, 1997, *87* (3), 363–382.

Gollin, Douglas, "Getting income shares right," *Journal of Political Economy*, 2002, *110*, 458–474.

Greenwood, Jeremy and Boyan Jovanovic, "The information-technology revolution and the stock market," *American Economic Review*, 1999, *89*, 116–122.

Helpman, Elhanan and Manuel Trajtenberg, "Diffusion of General Purpose Technologies," in Elhanan Helpman, ed., *General Purpose Technologies and Economic Growth*, Cambridge, MA: MIT Press, 1998, pp. 85–119.

Iwaisako, Tatsuro, "Optimal mix of R&D subsidy and patent protection in a heterogeneous-industry R&D-based growth model," *Journal of Economic Dynamics and Control*, 2023, *154*, 104723.

Jones, Charles I., "R&D Based Models of Economic Growth," *Journal of Political Economy*, 1995, *103*, 759–784.

_ , "The End of Economic Growth? Unintended Consequences of a Declining Population," *American Economic Review*, 2022, *112* (11), 3489–3527.

Jorgenson, Dale W., "Information technology and the U.S. economy," *American Economic Review*, 2001, *91*, 1–32.

_ **, Kevin J. Stiroh, Robert J. Gordon, and Daniel E. Sichel**, "Raising the Speed Limit: U.S. Economic Growth in the Information Age," *Brookings Papers on Economic Activity*, 2000, pp. 125–235.

Jovanovic, Boyan and Yaw Nyarko, "Stepping-Stone Mobility," *Carnegie-Rochester Conference Series on Public Policy*, 1997, *46*, 289–325.

Klein, Michael A. and Fuat Şener, "Product innovation, diffusion and endogenous growth," *Review of Economic Dynamics*, 2023, *48*, 178–201.

Laitner, John, "Structural Change and Economic Growth," *Review of Economic Studies*, 2000, *67*, 545–561.

Lin, Hwan C., "Creative Destruction and Optimal Patent Life in a Variety-Expanding Growth Model," *Southern Economic Journal*, 2013, *80* (1), 78–100.

_ , "Quality-Ladder Innovation and Overlapping Generations," *Economic Inquiry*, 2013, *51* (3), 1924–1942.

Mehlum, Halvor, Ragnar Torvik, and Simone Valente, "Demographic Change and the Sustainability of Economic Growth," *Journal of Economic Dynamics and Control*, 2020, *116*, 103916.

_ , _ , and _ , "Growth with Age-Dependent Preferences," *Journal of International Trade & Economic Development*, 2020, *29* (3), 320–344.

Microsoft, AI Economy Institute, "Global AI Adoption in 2025– A Widening Digital Divide," Technical Report, Microsoft 2025.

Pástor, Lŭboš and Pietro Veronesi, "Technological revolutions and stock prices," *American Economic review*, 2009, *99*, 1451–1483.

Romer, David, "The Solow Growth Model," in "Advanced Macroeconomics," McGraw-Hill, 2012, chapter 1, pp. 6–48.

Venkatesh, Viswanath, James Y. L. Thong, and Xin Xu, "Consumer Acceptance and Use of Information Technology: Extending the Unified Theory of Acceptance and Use of Technology," *MIS Quarterly*, 2012, *36* (1), 157–178.

A Household Optimality

A.1 Setup

Agent i born at t maximises $\log C^y_{i,t} + \beta \log C^o_{i,t+1}$ subject to (5)–(6).

A.2 Euler Equation

Substituting the budget constraints into the objective and differentiating with respect to $S_{i,t}$:

$$\frac{1}{C^y_{i,t}} = \frac{\beta(1+r_{t+1})}{C^o_{i,t+1}}.$$

Using $C^y_{i,t} = E^y_{i,t}/P^y_t$ and $C^o_{i,t+1} = E^o_{i,t+1}/P^o_{t+1}$:

$$\frac{C^o_{i,t+1}}{C^y_{i,t}} = \beta(1+r_{t+1})\frac{P^y_t}{P^o_{t+1}}. \tag{A.1}$$

A.3 Optimal Savings under Log Utility

From the period constraints and the Euler equation:

$$P^y_t C^y_{i,t}(1+\beta) = e_{i,t}, \quad \Rightarrow \quad P^y_t C^y_{i,t} = \frac{e_{i,t}}{1+\beta}, \quad S_{i,t} = \frac{\beta}{1+\beta} e_{i,t}. \tag{A.2}$$

The savings rate $\beta/(1+\beta)$ is independent of prices and the rental rate.

A.4 Price Index under Exponential Decay Weights

Substituting $\phi^a(s,\omega) = (1+\delta^a\omega)e^{-\theta^a s}$ into (4):

$$(P^a_t)^{1-\sigma} = \int_0^\infty e^{-(\sigma-1)\theta^a s}\left[(1+\delta^a)^{1/\sigma} p^{1-\sigma}_{\tau,1,t} + p^{1-\sigma}_{\tau,0,t}\right] d\tau. \tag{A.3}$$

Along the BGP with $p_{\tau,\omega,t} = \tilde{p}_\omega A^{1/(1-\alpha)}_{t-s}$ and $A_{t-s} = A_t(1+g^*)^{-s}$:

$$(P^a_t)^{1-\sigma} = A_t^{(\sigma-1)/(1-\alpha)} \int_0^\infty e^{-\lambda^a s}\left[(1+\delta^a)^{1/\sigma}\tilde{p}_1^{1-\sigma} + \tilde{p}_0^{1-\sigma}\right] ds, \tag{A.4}$$

where $\lambda^a = (\sigma-1)[\theta^a + g^*/(1-\alpha)]$. The integral converges for $\lambda^a > 0$, which holds since $\theta^a > 0$ and $g^* > 0$.

A.5 Aggregation over the Endowment Distribution

With i.i.d. endowments and log utility, aggregate savings are:

$$S_t = \int S_{i,t}\, di = \frac{\beta}{1+\beta}\bar{e}_t. \tag{A.5}$$

All aggregate quantities depend on the endowment distribution only through the mean $\bar{e}_t$.

B Demand Wave Dynamics

B.1 Proof of Proposition 1

Differentiating (23) with respect to $\tilde{\Delta}$:

$$\frac{\partial D_t^{\text{GPT}}}{\partial \tilde{\Delta}} = \frac{\bar{e}_t}{1+\beta}\frac{\partial \Phi_t^y}{\partial \tilde{\Delta}} + (1+\rho_t)S_{t-1}\frac{\partial \Phi_t^o}{\partial \tilde{\Delta}}. \tag{B.1}$$

From (22), $\partial \Phi_t^a / \partial \tilde{\Delta} > 0$ for both $a \in \{y, o\}$ since a higher $\tilde{\Delta}$ raises the GPT expenditure share of both cohorts. Since both terms are positive, $\partial D_t^{\text{GPT}} / \partial \tilde{\Delta} > 0$.

B.2 Composite Sufficient Statistic $\tilde{\Delta}$

The GPT expenditure share (22) can be written as:

$$\Phi_t^a = \frac{(1+\delta^a)^{(\sigma-1)/\sigma}/\lambda^a}{(1+\delta^a)^{(\sigma-1)/\sigma}/\lambda^a + 1/\lambda^a} = \frac{(1+\delta^a)^{(\sigma-1)/\sigma}}{(1+\delta^a)^{(\sigma-1)/\sigma} + 1}. \tag{B.2}$$

The composite statistic $\tilde{\Delta}$ arises from the difference in the demand-weighted integrals:

$$\tilde{\Delta} = \frac{(1+\delta^y)^{\sigma-1}}{\lambda^y} - \frac{(1+\delta^o)^{\sigma-1}}{\lambda^o} > 0, \tag{B.3}$$

where Assumption 4 ensures $\tilde{\Delta} > 0$.

B.3 Demand Wave Decomposition

The change in D_t^{GPT} at a cohort turnover decomposes as:

$$\Delta D_t^{\text{GPT}} = \underbrace{\frac{\bar{e}_t}{1+\beta}\Delta\Phi_t^y}_{\text{cohort-turnover wave}} + \underbrace{(1+\rho_t)S_{t-1}\Delta\Phi_t^o}_{\text{old-cohort response}} + \underbrace{\frac{\bar{e}_t}{1+\beta}\Phi_{t-1}^y \cdot \phi^y(0,1)}_{\text{vintage-novelty spike}}. \tag{B.4}$$

The third term is the vintage-novelty spike: when a new GPT variety enters at $\tau = t$, young agents immediately demand it at weight $\phi^y(0,1) = 1 + \delta^y$. This term is absent in the vintage-neutral model.

B.4 Stationarity of Φ_t^a on the BGP

On the BGP, the integrals $I_\omega^a = \int_0^\infty e^{-\lambda^a s} \tilde{p}_{s,\omega}^{1-\sigma} ds$ converge to constants (since $\tilde{p}_{s,\omega}$ is stationary). Therefore, Φ_t^a is time-invariant on the BGP, confirming that $\tilde{\Delta}$ is a valid sufficient statistic for the BGP analysis.

C Firm Optimality

C.1 Relative Capital Demand

From (24), the ratio of GPT to non-GPT capital demand is:

$$\frac{K_{\tau,1,t}}{K_{\tau,0,t}} = \left(\frac{A_{\tau,1,t}}{A_{\tau,0,t}}\right)^{1/(1-\alpha)} = \Gamma^{1/(1-\alpha)} \equiv \gamma > 1. \tag{C.1}$$

C.2 Sector-Level Aggregates

The GPT sector's share of aggregate capital:

$$s_t = \frac{\mu_t \gamma}{1 - \mu_t + \mu_t \gamma} \in (0,1), \tag{C.2}$$

which is strictly increasing and convex in μ_t (since $\gamma > 1$).

C.3 Proof of Proposition 2

From (29), $\Lambda(\tilde{\Delta}) = \Omega(\tilde{\Delta})\gamma - 1$. Differentiating $\Omega(\tilde{\Delta})$ with respect to $\tilde{\Delta}$:

$$\frac{\partial \Omega(\tilde{\Delta})}{\partial \tilde{\Delta}} > 0, \tag{C.3}$$

since $\Omega(\tilde{\Delta})$ is a weighted average of $(1+\delta^a)^{1/\sigma}$ with weights $1/\tilde{\lambda}^a$, and $\partial\Omega/\partial\delta^a > 0$ for both a. Hence $\partial\Lambda/\partial\tilde{\Delta} > 0$, and since $\rho(\tilde{\Delta})$ is a power function of $\Lambda(\tilde{\Delta})$ with positive exponent $(1-\alpha)/\alpha > 0$:

$$\frac{\partial \rho(\tilde{\Delta})}{\partial \tilde{\Delta}} = \frac{(1-\alpha)}{\alpha} \cdot \rho(\tilde{\Delta}) \cdot \frac{\partial \ln \Lambda(\tilde{\Delta})}{\partial \tilde{\Delta}} > 0. \tag{C.4}$$

C.4 Proof of Proposition 3 (R&D Complementarity)

From (34)–(35), the post/pre-adoption R&D ratio is:

$$\frac{R_{\tau,1,t}/A_{\tau,1,t}}{R_{\tau,0,t}/A_{\tau,0,t}} = \left(\frac{\xi_1 V_{\tau,1,t+1}}{\xi_0 V_{\tau,0,t+1}}\right)^{1/(1-\psi)}. \tag{C.5}$$

Since $\xi_1 = \zeta\xi_0 > \xi_0$ and $V_{\tau,1,t+1} = \gamma V_{\tau,0,t+1}$:

$$\frac{R_{\tau,1,t}/A_{\tau,1,t}}{R_{\tau,0,t}/A_{\tau,0,t}} = (\zeta\gamma)^{1/(1-\psi)} > 1, \tag{C.6}$$

since $\zeta > \Gamma > 1$ and $\gamma = \Gamma^{1/(1-\alpha)} > 1$.

D Market Clearing

D.1 Capital Market Clearing

With $\delta_k = 0$, capital fully depreciates within the period. Young agents' savings at $t-1$ become the capital stock at t:

$$K_t = S_{t-1} = \frac{\beta}{1+\beta}\bar{e}_{t-1}. \tag{D.1}$$

The equilibrium rental rate clears the capital market:

$$\rho_t = \alpha A_t^{1-\alpha} K_t^{\alpha-1} G(\mu_t)^{1-\alpha} = \alpha\left(\frac{A_t}{K_t}\right)^{1-\alpha} G(\mu_t)^{1-\alpha}, \tag{D.2}$$

where $G(\mu_t) = \mu_t\gamma + (1-\mu_t)$ is the composition index.

D.2 Goods Market Clearing

Total output equals total absorption:

$$Y_t = \frac{\bar{e}_t}{1+\beta} + \frac{\beta(1+\rho_t)\bar{e}_{t-1}}{1+\beta} + \frac{\beta\bar{e}_t}{1+\beta}. \tag{D.3}$$

The three terms are: young-cohort consumption, old-cohort consumption (from savings returns), and young-cohort savings (= next period's capital).

D.3 Equilibrium $\rho^* > \rho(\tilde{\Delta})$: Proof that Adoption is Interior

On the BGP, the equilibrium rental rate $\rho^* = \beta^{-1}(1+g^*) - 1$. The adoption threshold $\rho(\tilde{\Delta})$ from (33) depends on μ^* through $\kappa(\mu^*) = \kappa_0(\mu^*)^{-\varepsilon}$. The interior adoption condition $\mu^* \in (0,1)$ requires:

$$\rho(\tilde{\Delta})\big|_{\mu=0} < \rho^* < \rho(\tilde{\Delta})\big|_{\mu=1}, \tag{D.4}$$

which is satisfied under the parameter restrictions of Theorem 1.

E BGP Existence, Uniqueness, and Stability

E.1 Proof of Theorem 1

Existence. Define the BGP mapping $T : (\mu, A, K, \rho) \mapsto (\mu', A', K', \rho')$ by the system of equations (33), (40)–(41), and (36). A fixed point of T is a BGP. By the intermediate value theorem, a fixed point exists provided:

1. $\rho(\tilde{\Delta})|_{\mu=0} = 0$ (no adoption is not an equilibrium when $\Lambda(\tilde{\Delta}) > 0$).
2. $\rho(\tilde{\Delta})|_{\mu=1} < \infty$ (full adoption is feasible).
3. $\rho(\tilde{\Delta})$ is continuous and strictly increasing in μ.

All three conditions hold under Assumptions 3–4.

Uniqueness. Suppose two BGPs $\mu_1^* \neq \mu_2^*$ both satisfy $\rho^* = \rho(\tilde{\Delta})$. Since $\rho(\tilde{\Delta})$ is strictly increasing in μ (Proposition 2) and ρ^* is uniquely determined by (41), there is at most one μ^* satisfying $\rho^* = \rho(\tilde{\Delta})$.

Growth rate. From the R&D law of motion (13), constant growth requires $R_{\tau,\omega,t}/A_{\tau,\omega,t} = (g^*/\xi_\omega)^{1/\psi}$. Substituting into (34)–(35) and using the capital market clearing, the BGP growth rate is uniquely $g^* = g_a/(1-\alpha)$.

E.2 Stability of the BGP

Local stability. Linearising the transition map T around the BGP $(\mu^*, A^*, K^*, \rho^*)$, the Jacobian has eigenvalues:

$$\lambda_1 = 1 - \nu \in (0,1), \quad \lambda_2 = \frac{1}{1+g^*} \in (0,1), \quad \lambda_3 = \frac{\beta}{1+\beta} \in (0,1), \tag{E.1}$$

where ν is given by (50). All eigenvalues are in $(0,1)$, so the BGP is locally stable.

E.3 Proof of Proposition 4

From (25):

$$Y_t = \int_{-\infty}^{t} \pi_{\tau,\omega(\tau),t}\, d\tau = (1-\alpha)\left(\frac{\alpha}{\rho_t}\right)^{\alpha/(1-\alpha)} \int_{-\infty}^{t} A_{\tau,\omega(\tau),t}^{1/(1-\alpha)}\, d\tau.$$

Since $A_{\tau,1,t} = \Gamma A_{\tau,0,t}$:

$$\int A_{\tau,\omega,t}^{1/(1-\alpha)} d\tau = A_t^{1/(1-\alpha)} \left[\mu_t \Gamma^{1/(1-\alpha)} + (1-\mu_t)\right] = A_t^{1/(1-\alpha)} \left[1 + \mu_t(\gamma - 1)\right].$$

Therefore:

$$Y_t = (1-\alpha)\left(\frac{\alpha}{\rho_t}\right)^{\alpha/(1-\alpha)} A_t^{1/(1-\alpha)} \left[1 + \mu_t(\gamma - 1)\right]. \tag{E.3}$$

This is *linear* in μ_t holding ρ_t fixed. Superlinearity arises because ρ_t itself depends on μ_t through capital market clearing. Differentiating (37) with respect to μ_t:

$$\frac{\partial \rho_t}{\partial \mu_t} = (1-\alpha)\rho_t \cdot \frac{\partial \ln G(\mu_t)}{\partial \mu_t} = (1-\alpha)\rho_t \cdot \frac{\gamma - 1}{G(\mu_t)} > 0.$$

A higher μ_t raises ρ_t (more productive capital), which feeds back into higher profits. The second derivative:

$$\frac{\partial^2 Y_t}{\partial \mu_t^2} = 2(1-\alpha)^2 \rho_t \left(\frac{\alpha}{\rho_t}\right)^{\alpha/(1-\alpha)} A_t^{1/(1-\alpha)} \left(\frac{\gamma - 1}{G(\mu_t)}\right)^2 > 0.$$

E.4 Convergence Speed

From the linearised dynamics of μ_t:

$$\mu_{t+1} - \mu^* = (1-\nu)(\mu_t - \mu^*),$$

where $\nu = \varepsilon \kappa_0 (\mu^*)^{-\varepsilon-1} / (\partial \rho(\tilde{\Delta})/\partial \mu)|_{\mu^*}$. Since $\kappa_0 > 0$, $\varepsilon > 0$, and $\partial \rho(\tilde{\Delta})/\partial \mu > 0$, we have $\nu \in (0,1)$ and convergence follows geometrically at rate $1 - \nu$.

F Welfare Derivations

F.1 Social Planner's First-Order Conditions

The Lagrangian for the social planner's problem:

$$\mathcal{L} = \sum_{t=0}^{\infty} \beta^t \left[\log C_t^y + \beta \log C_{t+1}^o\right] + \sum_{t=0}^{\infty} \lambda_t \left[Y_t - C_t^y - C_t^o - K_{t+1} - \int R_{\tau,\omega,t}\, d\tau - \int \kappa(\mu_t)\, d\tau\right]. \tag{F.1}$$

FOC for C_t^y: $\beta^t / C_t^y = \lambda_t$.
FOC for C_{t+1}^o: $\beta^{t+1}\beta / C_{t+1}^o = \lambda_t$.
FOC for K_{t+1}: $\lambda_t = \lambda_{t+1}(1 + \rho_{t+1})$.
These yield the social Euler equation:

$$\frac{C_{t+1}^o}{C_t^y} = \beta(1 + \rho_{t+1}) \frac{P_t^y}{P_{t+1}^o}, \tag{F.2}$$

which coincides with the competitive Euler equation (18). Hence, the competitive equilibrium is constrained efficient, conditional on the adoption and R&D decisions.

F.2 Wedge Derivations

R&D wedge. The social FOC for $R_{\tau,\omega,t}$:

$$\lambda_t = \lambda_{t+1} \frac{\partial Y_{t+1}}{\partial A_{t+1}} \cdot \frac{\partial A_{t+1}}{\partial R_{\tau,\omega,t}} + \sum_{s>t+1} \lambda_s \frac{\partial Y_s}{\partial A_s} \frac{\partial A_s}{\partial R_{\tau,\omega,t}}. \tag{F.3}$$

The competitive firm's FOC only accounts for the first-period return. The wedge is the sum of future spillover terms.

Adoption wedge. The social FOC for μ_t:

$$\lambda_t \left[\frac{\partial Y_t}{\partial \mu_t} - \kappa(\mu_t) - \mu_t \kappa'(\mu_t)\right] = 0. \tag{F.4}$$

The competitive firm sets $V_{\tau,1,t} - V_{\tau,0,t} = \kappa(\mu_t)$, ignoring the externality term $-\mu_t \kappa'(\mu_t) > 0$.

F.3 Welfare Decomposition Proof

The welfare decomposition (72) follows from separating the resource constraint into consumption, adoption, and R&D components and using the envelope theorem

to evaluate each component at the optimum.

G Notation and Stationarity Summary

G.1 Notation Summary

Table 6: Notation summary.

Symbol	Definition	Notes
t	Calendar date	Discrete
τ	Vintage index	$\tau \leq t$
$s = t - \tau$	Vintage age	$s \geq 0$
$\omega \in \{0, 1\}$	GPT status	1=GPT, 0=non-GPT
$a \in \{y, o\}$	Agent age	y=young, o=old
$\phi^a(s, \omega)$	Demand weight	$(1 + \delta^a \omega) e^{-\theta^a s}$
δ^a	GPT taste premium	$\delta^y > \delta^o > 0$
θ^a	Vintage decay rate	$\theta^y > \theta^o > 0$
$\delta = \delta^y - \delta^o$	GPT taste gap	> 0
$\Theta = \theta^y - \theta^o$	Vintage-discount gap	> 0
λ^a	Effective decay rate	$(\sigma - 1)[\theta^a + g^*/(1 - \alpha)]$
$\tilde{\Delta}$	Composite sufficient statistic	$(1 + \delta^y)^{\sigma-1}/\lambda^y - (1 + \delta^o)^{\sigma-1}/\lambda^o$
μ_t	Adoption rate	$\in [0, 1]$
$\gamma = \Gamma^{1/(1-\alpha)}$	GPT capital multiplier	> 1
$\kappa(\mu) = \kappa_0 \mu^{-\varepsilon}$	Adoption cost	Decreasing in μ
ρ_t	Rental rate (= net return)	$\delta_k = 0$
ρ^*	BGP rental rate	$\beta^{-1}(1 + g^*) - 1$
$\rho(\tilde{\Delta})$	Adoption threshold	Increasing in $\tilde{\Delta}$
g^*	BGP growth rate	$g_a/(1 - \alpha)$
$\Lambda(\tilde{\Delta})$	Net GPT revenue multiplier	$\Omega(\tilde{\Delta})\gamma - 1 > 0$
$\Omega(\tilde{\Delta})$	Demand-weighted revenue ratio	Defined in (28)
s_t	GPT capital share	$\mu_t\gamma/(1 - \mu_t + \mu_t\gamma)$
t_D^*	Demand wave peak	$\arg\max D_t^{\text{GPT}}$
ν	Richards asymmetry parameter	> 0
t_S^*	S-curve inflection	$\arg\max \dot{\mu}_t$
$\tau^* = t_S^* - t_D^*$	Lead-lag offset	> 0

G.2 Stationarity of Key Variables on the BGP

Table 7: Stationarity of key variables on the BGP.

Variable	BGP growth rate	Notes
$A_{\tau,\omega,t}$	$g^* = g_a/(1-\alpha)$	Common across sectors
K_t	g^*	Capital accumulation
Y_t	g^*	Output
$\bar{e}_t$	g^*	Mean endowment
ρ_t	0 (stationary)	Rental rate
μ_t	0 (stationary)	Adoption rate
Φ_t^a	0 (stationary)	GPT expenditure shares
$\tilde{\Delta}$	0 (stationary)	Composite statistic
$V_{\tau,\omega,t}$	g^*	Firm values
$\pi_{\tau,\omega,t}$	g^*	Firm profits
s_t	0 (stationary)	Capital shares

H Data Sources and Reproducibility

Data. All data sources are publicly available without restricted access:

1. **Ken French Data Library** — 49 Industry Portfolios (annual returns, 1927–2025) and Fama-French 3-Factor Model (monthly, 1926–2026): *Files used:* `49_Industry_Portfolios.csv` (industry 35 = Computers), `F-F _Research _Data_Factors.csv` (market excess return).

2. **NBER Consumer Expenditure Survey (CEX/CBO) Extracts** — Fixed-format family and member records (1990–2003): *Files used:* `ffile{yy}{q}` (family expenditures, 109 CBO variables), `mfile{yy}{q}` (member demographics, age at positions 8–10).

3. **BLS Consumer Expenditure Tables** — Published cross-tabulations by age of reference person: available at `bls.gov/cex/tables.htm` (BLS Table 1101 snapshot for the 1990s).

Replication code. The calibration is fully reproducible via three Python scripts:

- `phase3_richards_scurve.py` — Richards curve calibration (7 sections): solves for ν, computes adoption path, moments, equity premium, and sensitivity analysis. Outputs: `phase3_richards_results.json`.
- `cex_data/parse_cex_micro.py` — NBER CEX micro-data parser: extracts age from member records, merges with family expenditures, computes Young/Old ratios. Outputs: `cex_micro_results.json`.