University of Warwick

Research in Applied Economics

Maximilien Delaporte
u2126980

8th of May 2025

Word Count: 4,996[1]

[1] **Including**: Main Body, Footnotes, Headings, Subheadings, and Equations (5 words per line)
Excluding: Title Page, Abstract, Acknowledgements, Table of Contents, In-Text References, Tables & Figures, Appendix, and Bibliography

Nannies, Machines, and Population Growth:
How Could the Automation of Childcare Affect Economic Growth?

Abstract

This dissertation investigates the potential for artificial intelligence (AI) to impact long-run economic growth through its effect on fertility. While recent Semi-Endogenous Growth (SEG) frameworks have explored AI-driven automation in production and idea creation, little attention has yet been given to how automation may alter household decisions by automating parenting tasks. We develop a task-based Semi-Endogenous Growth model in which fertility is endogenously determined and influenced by the level of AI-driven automation in childcare. The model integrates insights from Zeira (1998), Doepke et al. (2023), and Jones (2022), allowing automation to alter the cost distribution of childrearing and impact household fertility decisions.

A closed-form solution for fertility is derived under time and budget constraints. Preliminary results suggest that automation could significantly raise fertility in certain settings, providing a novel pathway through which policy can aim to mitigate the global decline in fertility.

This framework offers a new theoretical channel for understanding automation's role in sustaining growth and suggests that household-level automation has the potential to become an increasingly relevant component of demographic policy discussions.

Acknowledgements

I would like to thank my dissertation supervisor, Dr. Cholwoo Kim, for challenging me as well as providing continuous support and insightful feedback throughout my research project. I would also like to thank my professor, Dr. Federico Rossi, whose lectures on semi-endogenous growth theory inspired this paper, and whose help was crucial in the development of this model.

Contents

1. Introduction

Research Motivation, Question, and Contribution

Modern growth theory sees long-run economic growth as fundamentally reliant on the continuous production of new technologies[2], which increase the productivity of workers. Within the Semi-Endogenous Growth (SEG) framework specifically, new technologies become harder and harder to find as the economy's technological level increases.

Population growth then becomes a crucial component of long-run economic growth as it provides a way to mitigate the decreasing returns to innovation by supplying more idea-generating agents. Jones (2022) implies that (currently) population growth is our only source of long-term economic growth, despite accounting only for roughly 20% of observed economic growth of the past +50 years.[3]

Concerningly, however, global fertility rates have been falling steadily; the UN's projections (which we plot on Figure 1) now predict that global population will peak around 10.3 Billion in the mid-2080s (United Nations, 2024), posing a significant structural threat to the only long-term driver of economic growth, according to SEG theory.[4]

Figure 1: *UN Population Growth Estimations and Projections (1950-2100)*

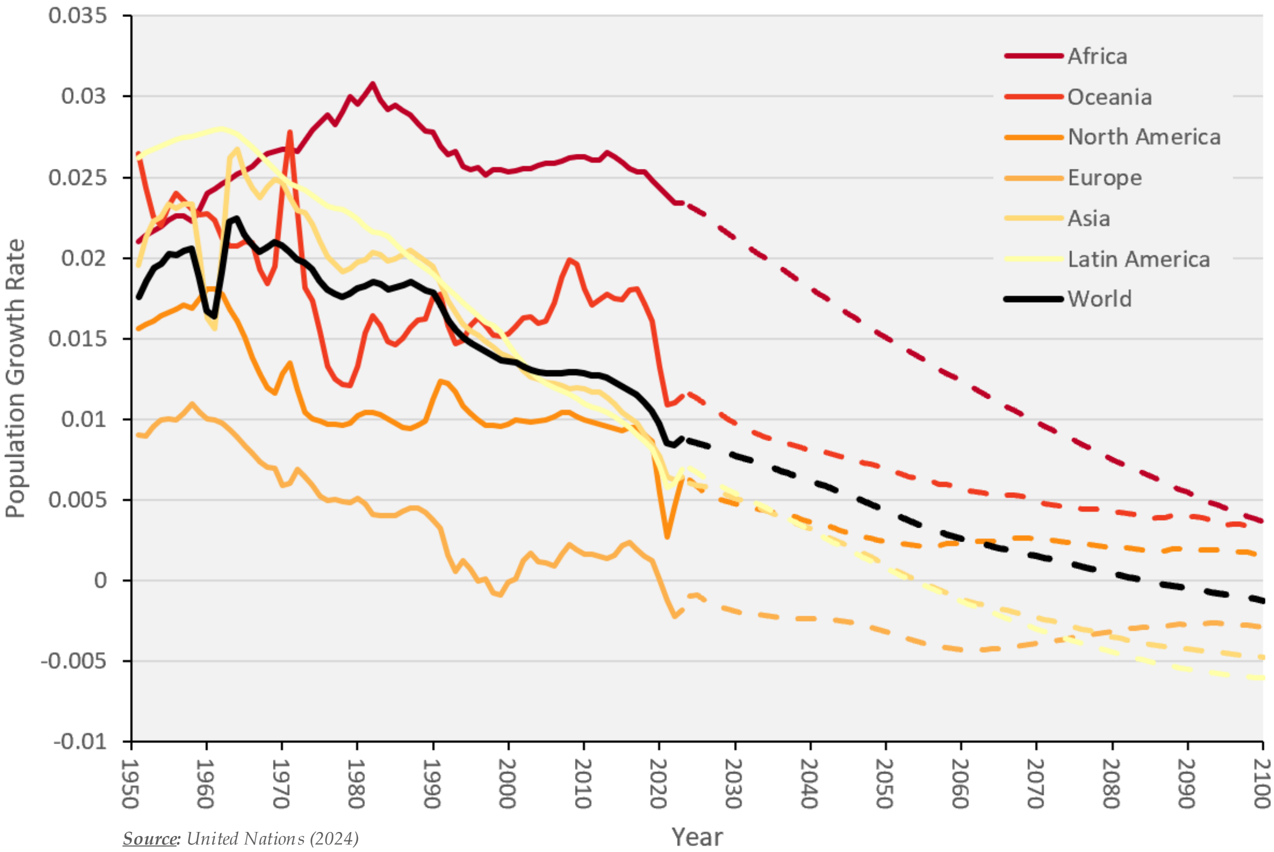

Source: *United Nations (2024)*

[2] Also referred to as ideas, knowledge, or innovation.

[3] He explains the remaining 80% as ephemeral level effects; bound to gradually stop driving global economic growth.

[4] As illustrated in Jones (2020, 2022) with the situation dubbed 'the empty planet scenario' – where the fall of population growth leads to stagnating knowledge and living standards for a population that vanishes.

Doepke et al. (2023) propose a model for endogenous fertility which explains its gradual decrease as resulting from two main underlying trends:

(1) Higher female wages which, in turn, increase the relative time opportunity-cost of childrearing, making it a less attractive option.
(2) A decreasing societal preference for large families, possibly influenced by Malthusian fears of overpopulation.

Turning to a seemingly different but closely related matter, technological change – specifically the advent of general-purpose artificial intelligence (AI) – is reshaping sectors traditionally resilient to automation, including caregiving and education. Evidence from Japan's deployment of robotic eldercare analysed by Broadbent et al. (2009), the efficiency of AI-assisted learning platforms evaluated by Luo and Hsiao-chin (2023) – even the emergence of the satirical term 'iPad Kid'[5] – all signals point to technology playing an increasingly important role in childrearing. This shift, emphasized by the development of AI, suggests that automation – a biproduct of technological advancement – could increasingly substitute for parenting labour, potentially impacting fertility rates by altering household dynamics.

Yet, no existing economic framework explicitly connects technological change to household-level fertility choices, signalling an overlooked channel which this paper will focus on by addressing the following research question:

Given the new implications of AI for automation, how might applying it to childcare affect the household fertility dynamics that govern long-run economic growth? Could the automation of childcare help mitigate the causes of Jones' (2020) empty planet scenario?

To answer this, we develop a theoretical framework rooted in the SEG setting, where fertility is endogenized as the outcome of a utility maximisation problem à la Doepke et al. (2023). Households balance the consumption of goods and the number of children under both time and budget constraints. We structure the cost of childrearing using a task-based approach inspired by Zeira (1998), which enables its gradual automation.

A key innovation of the model is the introduction of an automation boundary, which represents the level of childcare automation, endogenously determined by the economy's technological level. As the level of technology is raised, the automation boundary expands, progressively reducing the time burden of childrearing and influencing household fertility choices.

By creating a link between the endogenous creation of technology from Jones' (2022) SEG framework to the level of childrearing automation, the model offers a novel mechanism, circularly linking technological progress to the demographic outcomes that govern it. In so doing, it contributes not only to the literature on automation and economic growth but also to the broader question of how technology might fundamentally alter the demographic foundations of idea production, labour supply, and long-run economic growth.

[5] Coined in popular culture to describe children given excessive and arguably harmful screen time for parental convenience (Parents, 2024).

2. Literature Review

This section reviews the strands of literature that inform the construction of the dissertation's theoretical model. It is divided into three parts: the SEG framework and its relationship to population, the integration of AI-driven automation within economic growth models, and recent attempts to link AI with fertility dynamics. Throughout, we highlight key contributions and identify where existing models leave space for the innovation proposed in this dissertation.

2.1. Semi-Endogenous Growth, Population, and AI Automation

The SEG framework, formalized by Jones (1995, 2022), asserts that sustained long-run growth requires ongoing idea production, which faces diminishing returns over time. Population growth is critical in this framework, as it supplies the increasing number of researchers necessary to counteract these decreasing returns. However, declining fertility rates globally (United Nations, 2024) threaten this mechanism, creating a vulnerability in traditional SEG models.

Attempts to address this vulnerability have incorporated AI-driven automation into SEG settings. Aghion et al. (2019) propose that AI could replace human labour not only in final goods production but also in idea creation, mitigating the reliance on population growth. However, their model introduces a significant limitation: under plausible parameter values, it predicts singularities where output grows infinitely in finite time, an empirically and theoretically unrealistic result. Jones (2022) notes this flaw and suggests that although AI may substitute for human inputs, its effects on demographic dynamics remain underexplored.

A related contribution by Gries and Naudé (2020) circumvents the singularity issue by treating idea creation as exogenous and focusing instead on endogenous technology adoption. While their approach stabilizes the growth path, it decouples population size from innovation, limiting its applicability for models concerned with demographic feedbacks. Thus, while AI's potential to sustain growth has been recognized, the interaction between AI-driven automation and fertility decisions remains largely absent from the current SEG literature.

2.2. Task-Based Models of Automation and Household Dynamics

The task-based approach to modelling automation, originating with Zeira (1998) and expanded by Acemoglu and Restrepo (2018), emphasizes heterogeneity in automation's impact across sectors and tasks. Zeira's partial equilibrium model introduced the concept of automation displacing labour in task-specific ways, while Acemoglu and Restrepo generalized this into a dynamic, general equilibrium setting that integrates new task creation alongside displacement. An automation frontier separates tasks performed by capital from those performed by labour, based on a cost cutoff. Although these frameworks richly capture market-sector dynamics, they have not been applied to childcare.

Critically, Acemoglu and Restrepo's frontier advances when high wages make automation profitable, so population growth - which tends to depress wages - *reduces* the incentive to automate, opposite to semi-endogenous-growth logic where more people → more researchers → faster innovation.

By contrast, household adoption of AI childcare tools hinges on technological capability (speech-recognition accuracy, safety standards) rather than parental wages. This motivates our use of a tech-pushed automation boundary: growing purely as a function of technology, featuring an explicit capability threshold A_{crit}, and capping out at a ceiling ΔN of automatable tasks. Transposing and adapting the task-frontier concept in this way captures the direct link between tech-readiness and a reduced time costs of parenting - an essential channel for the fertility feedback we want our model to capture.

Incorporating these insights into growth models has enhanced the realism of automation effects on labour markets and aggregate output. However, existing applications of task-based automation – including those in Aghion et al. (2019) and Jones (2022) – largely focus on production and idea creation sectors. Little attention has been paid to how task automation could penetrate non-market sectors such as household labour, specifically parenting tasks that heavily influence fertility decisions.

This omission is significant. The same mechanisms that make non-routine cognitive tasks automatable in production settings suggest that parts of childcare and early education could similarly be automated (Broadbent et al., 2009; Luo and Hsiao-chin, 2023). Current growth models thus overlook a potentially important channel through which AI can indirectly but materially alter macroeconomic outcomes by affecting household-level decision-making.

2.3. AI and Fertility: Indirect and Direct Mechanisms

The sparse literature on AI and fertility generally treats AI's influence as operating through labour market adjustments. Wei and Xie (2022) model AI-induced changes in wages and working hours as affecting fertility indirectly by altering the opportunity cost of childbearing. While valuable, this approach assumes that AI does not alter the production function of raising children itself, only the broader economic environment in which fertility decisions are made.

By contrast, the model proposed in this dissertation explores a direct channel: AI-driven automation of childcare tasks reduces the time and financial costs of childrearing, modifying the household utility-maximization problem at its core. This represents a substantive departure from existing frameworks, which generally treat the time and cost parameters in fertility models as technologically static (Doepke et al., 2023).

Thus, while current models recognize AI's transformative economic potential, they underestimate its ability to alter demographic foundations directly. Addressing this gap by endogenizing task automation within the fertility decision is essential to understanding the full macroeconomic consequences of recent technological breakthroughs.

3. The Model

3.1. Overview and Intuition

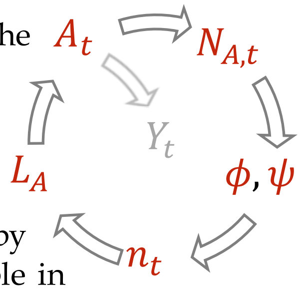

The intuition of the model is as follows: the level of technology A_t (set by the SEG model) sets the level of childcare automation $N_{A,t}$. This in turn determines how many childrearing tasks incur a monetary ψ (automated) rather than a time ϕ (performed by parents) cost. The distribution of these costs directly impact households' fertility choice n, which drives the population growth needed for technological development, and thereby economic growth. A summary of the model's setup equations is available in Appendix (A.2.)

3.2. Growth Dynamics

The production side of the economy follows Jones' (2022) standard SEG model, composed of two production sectors: final goods and ideas.

Final Good Output $\boldsymbol{Y_t}$ is produced according to the following Cobb-Douglas function:

$$\boldsymbol{Y_t = A_t^{\sigma_y} L_{Y,t}} \qquad (1)$$

Where $\boldsymbol{A_t}$ represents the stock of ideas (or level of technology) at time t, $\boldsymbol{L_{Y,t}}$ is the amount of labour allocated to final good production, and $\boldsymbol{\sigma_y}$ is the elasticity of output with respect to the stock of knowledge.

The growth rate of technology $\boldsymbol{\dot{A}_t / A_t}$[6] is determined by equation (2) in which $\boldsymbol{z}$ represents the productivity of researchers, $\boldsymbol{L_{A,t}}$ is the number of researchers in the economy, and $\boldsymbol{\beta} > 0$ captures the decreasing returns to technological advancement.

$$\boldsymbol{\frac{\dot{A}_t}{A_t} = z L_{A,t} A_t^{-\beta}} \qquad (2)$$

As A_t grows, idea production becomes increasingly difficult, necessitating a larger research workforce to sustain the same rate of technological growth.

3.3. Population Dynamics

The economy's total workforce $\boldsymbol{L_t}$ is exogenously split between production and research by parameter $\boldsymbol{l} \in [0,1]$:

$$\boldsymbol{L_t = L_{A,t} + L_{Y,t}} \qquad \boldsymbol{\frac{L_{A,t}}{L_{Y,t}} = l} \qquad (3,4)$$

[6] Where the dot notation $\dot{A}_t$ denotes the time derivative of A_t.

While population growth is determined endogenously through household fertility decisions. Specifically, the **growth rate of the total workforce $\dot{L}_t/L_t$** during period t is given by:

$$\frac{\dot{L}_t}{L_t} = n_t - m \quad (5)$$

Where m represents an exogenous per-household mortality rate parameter and n_t denotes the economy's number of children per household at time t, endogenously determined by a household-level utility maximisation problem (section 3.5.). Essentially, equation (5) serves as the link which ties the outcome of our fertility decision model (n_t) to population growth within the SEG model.

3.4. Automation of Childcare

Before setting the household's maximisation problem, we define the mechanism through which the technological output of the economy $\dot{A}_t$ can influence fertility decisions by increasing the automation of childcare as the economy's technology level A_t rises. This is the principal contribution of this model to the literature, as it creates a novel link between macro-economic outcomes (in this case technological growth) and demographic tendencies. It does so through the creation of the **automation boundary $N_{A,t}$** – endogenously determined as a function of A_t – which represents the level of childcare automation.

Drawing from the task-based approach (Zeira 1998, Acemoglu and Restrepo 2018), our model sees childrearing as a collection of parenting tasks. Each period, parents **must** complete an exogenous number N of tasks per child in their household. The per-period effort E_t required to raise a child is therefore given by equation (6), where X represents the cost of task $i \in [1, N]$.

$$E_t = \sum_{i=1}^{N} X_{i,t} \quad (7)$$

When a task becomes automated it goes from incurring a **time cost ϕ** to a **monetary cost ψ** such that:

$$X_{i,t} = \begin{cases} w\phi & \text{if } i \in (N_{A,t}, N] \quad \text{(manual task)} \\ \psi & \text{if } i \in [1, N_{A,t}] \quad \text{(Automated task)} \end{cases} \quad (8)$$

Equation (7) can then be rewritten as:

$$E_t = \sum_{i=1}^{N_{A,t}} \psi + \sum_{i=N_{A,t}+1}^{N} w\phi = \psi N_{A,t} + w\phi(N - N_{A,t}) \quad (9)$$

Childrearing efforts are then endogenously split between time and monetary costs according to the automation boundary $N_{A,t}$, whose laws of motion are defined by the following equation, inspired by Acemoglu and Restrepo's "automation frontier", but developed as an original contribution of this model:

$$\dot{N}_{A,t} = A_t^{\sigma_n}\left(1 - \frac{N_{A,t}}{\Delta N}\right)\left(1 - e^{\kappa(A_{\text{crit}} - A_t)}\right) \quad (6)$$

Each period, the amount of newly automated tasks $\dot{N}_{A,t}$ is determined by three factors, each representing an underlying assumption of how our model sees the potential automation of childcare:

- $\boldsymbol{A_t^{\sigma_n}}$: The level of automation of childcare tasks depends on the economy's technology level (raised to an elasticity parameter $\boldsymbol{\sigma_n}$). A_t is the only endogenous determinant of $N_{A,t}$.

- $\left(\mathbf{1} - \frac{N_{A,t}}{\Delta N}\right)$: Some childrearing tasks cannot be automated. Certain aspects of parenting – such as emotional bonding, moral guidance, or unstructured play - are inherently resistant to automation. Regardless of how advanced technology becomes, certain tasks will always remain a parent's responsibility.
  The model therefore imposes a ceiling on childcare automation by scaling the increase in $N_{A,t}$ by $1 - \frac{N_{A,t}}{\Delta \mathrm{N}}$, where parameter $\boldsymbol{\Delta} \in [0,1)$ represents the maximum portion of tasks that can be automated. As $N_{A,t}$ approaches ΔN, automation gradually slows down to a stop.

- $\left(\mathbf{1} - e^{\kappa(A_{\text{crit}} - A_t)}\right)$: The automation of childcare can only begin once the economy surpasses an exogenously determined technology level $\boldsymbol{A_{\text{crit}}}$. The underlying assumption here is that a significant level of technology must be reached before automation reaches the childcare sector.
  This mechanic is modelled through the exponential smoothing term $1 - e^{\kappa(A_{\text{crit}} - A_t)}$, where the $\boldsymbol{\kappa}$ parameter controls how gradually automation ramps up once A_t exceeds A_{crit}.

These dynamics are simulated and illustrated in Section 4.1 (Figure 2)

3.5. Household Childbearing Decisions

As previously outlined (Section 3.3., Equation 5), our model's population growth rate is obtained by subtracting the exogenous mortality rate m from the per-household fertility rate n_t. This formulation enables population growth to depend on the number of children per household n_t, endogenously determined by a household utility maximisation framework à la Doepke et al. (2023). To close our model, we will incorporate our Zeira (1998) inspired automation dynamics, whereby technological progress influences the cost distribution of our household utility maximization problem, thus inter-linking technological progress, automation, population dynamics, and economic growth.

3.5.1. Household Preferences

As in Doepke et al. (2023), households derive utility from both consumption $\boldsymbol{c}$ and number of children $\boldsymbol{n}$. Preferences are represented by a separable log-utility function:

$$u(c, n) = \log(c) + \delta \log(n) \tag{10}$$

Where $\boldsymbol{\delta} > 0$ captures the relative importance of fertility compared to consumption. Households face both time and financial constraints when deciding how many children to have. Crucially, the model allows for technological progress to reduce these costs via automation of childcare tasks.

3.5.2. *Time and Budget Constraints*

Both agents of the household spend all their time either working or performing childrearing tasks. The time portion $\boldsymbol{\lambda_t}$ of period t that the household spends working is therefore equal to 1 minus the time spent taking care of its children:

$$\lambda_t = 1 - n_t \phi\left(N - N_{A,t}\right) \tag{11}$$

If the household has no children, it will spend all of period t working which will yield wage $\boldsymbol{w}$. For simplicity, we assume households do not either save or invest. The capital they earn during period t is entirely spent either on their children or on consumption.

$$w\lambda_t = c_t + n_t \psi N_{A,t} \tag{12}$$

3.5.3. *Solving for $\boldsymbol{n_t}$*

We substitute (11) into (12) and set up the household's utility maximisation problem:

$$\max_{c_t, n_t}: \log(c_t) + \delta \log(n_t) \quad s.t. \quad w\left(1 - n_t\left(N - N_{A,t}\right)\phi\right) = c_t + n_t N_{A,t} \psi \tag{13}$$

By substituting the budget constraint into the objective function and taking the first-order condition, we solve for a closed-form solution for optimal fertility (Appendix A.2.).

$$n_t = \frac{\delta}{1 + \delta} \cdot \frac{w}{w\phi\left(N - N_{A,t}\right) + \psi N_{A,t}} \tag{14}$$

3.6. Balanced Growth Path (BGP)

The BGP is a situation in which the model's major variables all grow at constant rates. To find an expression for those rates, we first calculate the growth rate $\boldsymbol{g_y}$ of output per capita $\boldsymbol{y}$.

$$y_t = \frac{Y_t}{L_t} = \frac{A_t^{\sigma_y} L_{Y,t}}{L_t} \tag{15}$$

Recalling equations (3) and (4):

$$L_{Y,t} = (1 - l)L_t \quad \Rightarrow \quad y_t = A_t^{\sigma_y}(1 - l) \tag{16}$$

Which yields:

$$g_y = \sigma_y g_A \tag{17}$$

3.6.1. *Long-Run Technological Growth* (g_A)

Recalling the technological growth equation (2) as well as equations (3) and (4), $\frac{\dot{A}_t}{A_t}$ can be expressed as:

$$\frac{\dot{A}_t}{A_t} = zlL_tA_t^{-\beta} \tag{18}$$

Which, along the BGP, yields the (constant) technological growth rate $\boldsymbol{g_A}$.

$$g_A = \frac{g_L}{\beta} \tag{19}$$

3.6.2. *Long-Run Population Growth* (g_L)

Recalling equations (10) and (12) we know that – as a function of A_t – $N_{A,t}$ converges either to zero (if $A_t < A_{\text{crit}}$) or to its saturation point ΔN. Therefore, along the BGP, n_t converges to:

$$n^* = \frac{\delta}{1+\delta} \cdot \frac{w}{w\phi(N-\Delta N)+\psi\Delta N} \tag{20}$$

The population growth rate along the BGP is therefore given by:

$$g_L = n^* - m \tag{21}$$

3.6.3. *Long-Run Output Growth* (g_y)

Substituting (19) and (21) into (17) gives us the constant growth rate of output g_y along the BGP:

$$g_y = \sigma_y \frac{n^* - m}{\beta} \tag{22}$$

In keeping with standard SEG theory, the long-term growth rate of output is dependent on fertility's relative strength to the decreasing rate of idea production β. The nuance our model brings, however, is that it treats population growth as an endogenous variable – rather than a constant rate – whose level depends on a multitude of underlying assumptions about automation's potential, in light of the new implications that AI has brought to the debate.

4. Quantitative Analysis

4.1. Simulated Dynamics of $N_{A,t}$

As discussed in section 4.4., the automation boundary is endogenously determined as a function of A_t (Equation 9). To better illustrate the dynamics of our novel automation boundary, we run a simulation[7] where we plot the evolution of $N_{A,t}$ in function of A_t, ceteris paribus.

Figure 2: Automation Boundary $N_{A,t}$ in Function of Technology Level A_t

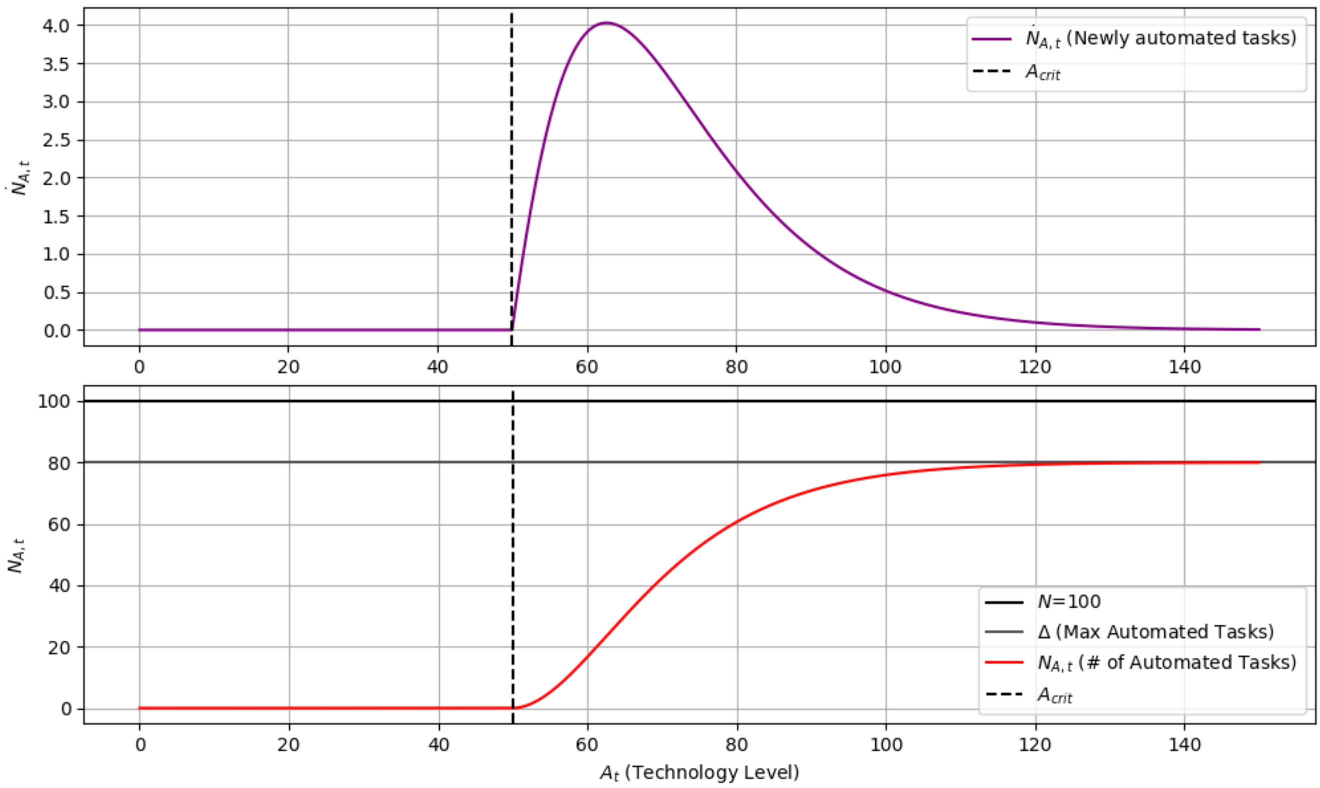

Figure 2 above shows the output of that simulation, from which we observe that plotting $N_{A,t}$ against A_t yields a sigmoidal[8], or s-shaped, curve which only departs from 0 once $A_t \geq A_{\text{crit}}$. Inputting different values for κ in the algorithm will change the rate at which childrearing tasks are replaced by automation (slope of the red line on Figure 2). Changing A_{crit} changes where along the x axis $N_{A,t}$ starts accumulating, and changing Δ sets the upper bound for our automation boundary $N_{A,t}$.

[7] The code scripts for all simulations are included in Appendix (A.4.)

[8] But not logistic since the acceleration and slowdown of $N_{A,t}$ in function of A_t is controlled by 2 different terms, allowing them to be asymmetric.

4.2. n_t Comparative Statics

We recall the fertility solution (Equation 14), run comparative statics analysis on it, and conduct ceteris paribus simulations to analyse the behaviour of our fertility rate over changes in all its determining factors.

$$n_t = \frac{\delta}{1+\delta} \cdot \frac{w}{w\phi(N - N_{A,t}) + \psi N_{A,t}} \tag{23}$$

4.2.1. *Preference for Childrearing ($\boldsymbol{\delta}$)*

$$\frac{\partial n_t}{\partial \delta} = \frac{1}{(1+\delta)^2} \cdot \frac{w}{w\phi(N - N_{A,t}) + \psi N_{A,t}} \tag{24}$$

$$\Rightarrow \frac{\partial n_t}{\partial \delta} > 0 \tag{25}$$

As in Doepke et al. (2023), an increase in households' preference for childrearing, ceteris paribus, unambiguously results in an increase in the fertility rate. The slope at which it increases is shown in the plotted simulation on Figure 3.

Figure 3: Plotted Simulation of n_t over δ

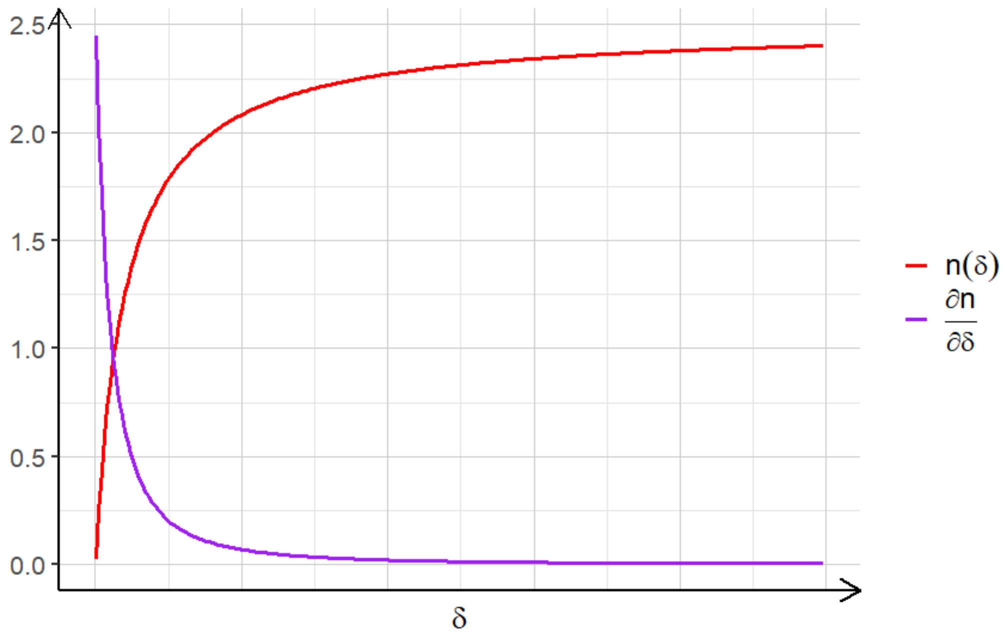

4.2.2. *Average Time-Cost of Childrearing Tasks ($\boldsymbol{\phi}$)*

$$\frac{\partial n_t}{\partial \phi} = \frac{\delta}{1+\delta} \cdot \frac{w^2(N - N_{A,t})}{\left[w\phi(N - N_{A,t}) + \psi N_{A,t}\right]^2} \tag{26}$$

$$\Rightarrow \frac{\partial n_t}{\partial \phi} \text{ is } \begin{cases} < 0 \text{ if: } & N_{A,t} < N \\ = 0 \text{ if: } & N_{A,t} = N \end{cases} \tag{27}$$

Equations (26) and (27) show that as the average time-cost ϕ rises, fertility n_t declines whenever $N_{A,t} < N$, with the effect disappearing when automation is complete ($N_{A,t} = N$). The magnitude of ϕ's impact decreases gradually as $N_{A,t}$ approaches N, meaning the sensitivity is proportional to the share of unautomated tasks. The simulations below (Figures 4 and 5) illustrate these dynamics.

Figure 4: Plotted Simulation of n_t over ϕ ($N_{A,t} = 0$)

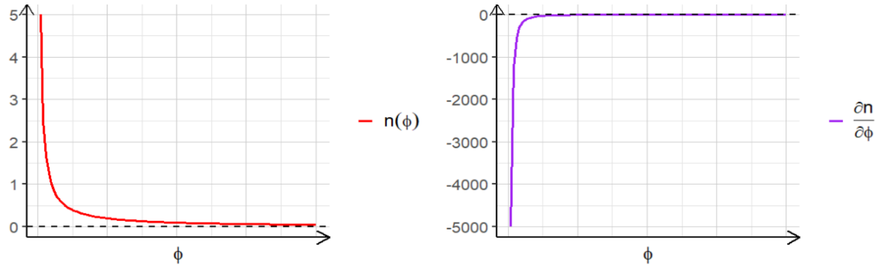

Figure 5: Plotted Simulation of n_t over ϕ ($N_{A,t} = N$)

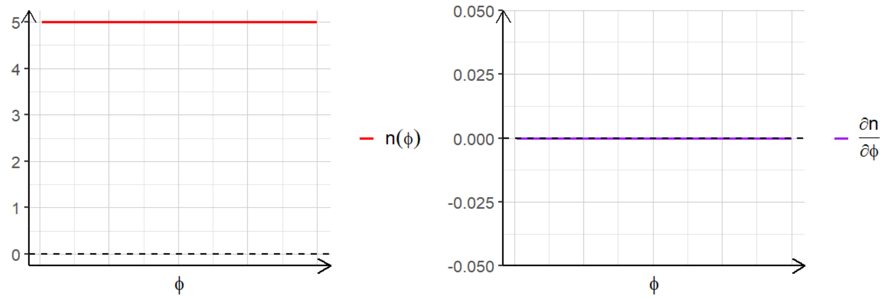

4.2.3. *Average Monetary Cost of Automated Tasks (ψ)*

$$\frac{\partial n_t}{\partial \psi} = \frac{\delta}{1+\delta} \cdot \frac{wN_{A,t}}{\left[w\phi(N - N_{A,t}) + \psi N_{A,t}\right]^2} \tag{28}$$

$$\Rightarrow \frac{\partial n_t}{\partial \psi} \text{ is } \begin{cases} < 0 & \text{if: } N_{A,t} > 0 \\ = 0 & \text{if: } N_{A,t} = 0 \end{cases} \tag{29}$$

Similarly to ϕ, raising ψ unambiguously lowers the fertility rate n_t. However, the magnitude of this effect is inversely scaled to ϕ: as the number of automated tasks increase, the effect of raising ϕ on fertility diminishes while that of raising ψ increase – as shown in Figures 6 and 7.

Figure 6: Plotted Simulation of n_t over ψ ($N_{A,t} = 0$)

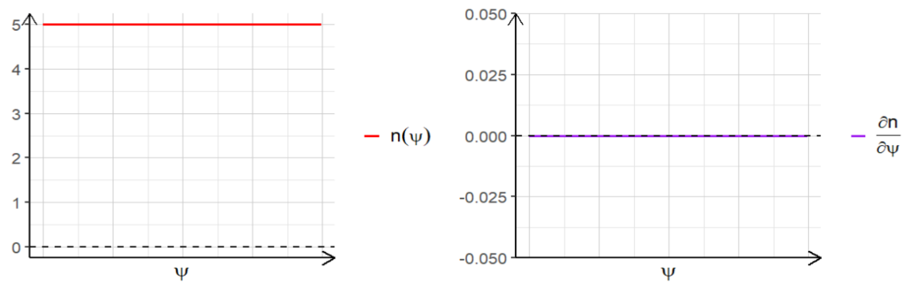

Figure 7: *Plotted Simulation of n_t over ψ ($N_{A,t} = N$)*

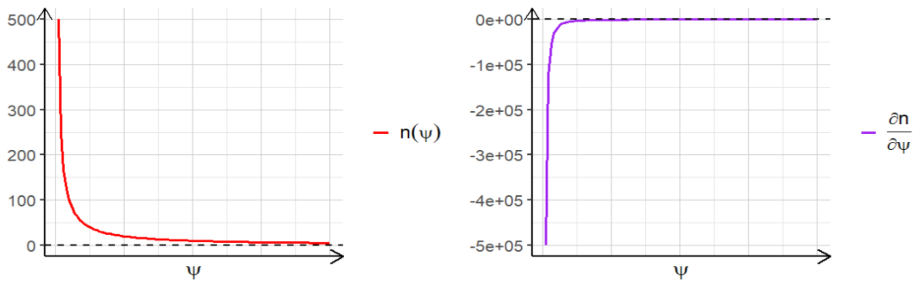

4.2.4. *Wages (w)*

$$\frac{\partial n_t}{\partial w} = \frac{\delta}{1+\delta} \cdot \frac{\psi N_{A,t}}{\left[w\phi\left(N - N_{A,t}\right) + \psi N_{A,t}\right]^2} \tag{30}$$

$$\Rightarrow \frac{\partial n_t}{\partial w} \text{ is } \begin{cases} > 0 \text{ if: } \psi N_{A,t} > 0 \\ = 0 \text{ if: } \quad N_{A,t} = 0 \end{cases} \tag{31}$$

Equations (30) and (31) show that as the wages increase, fertility rises proportionally to the level of $N_{A,t}$, with no effect when automation is absent ($N_{A,t} = 0$). As long as $N_{A,t} \neq N$, the model exhibits decreasing returns to wage increases, and n_t converges to a constant as $w \rightarrow +\infty$, since the unautomated tasks impose a hard ceiling on fertility gains. By contrast, if full automation is achieved ($N_{A,t} = N$), fertility increases linearly with w. Since that situation implies there are no unautomated childcare tasks, the time cost of childrearing is null and household fertility choices become solely dependent on how many children a household can afford. The simulations below (Figures 8, 9, and 10) illustrate these dynamics.

Figure 8: *Plotted Simulation of n_t over w ; ($N_{A,t} = 0$)*

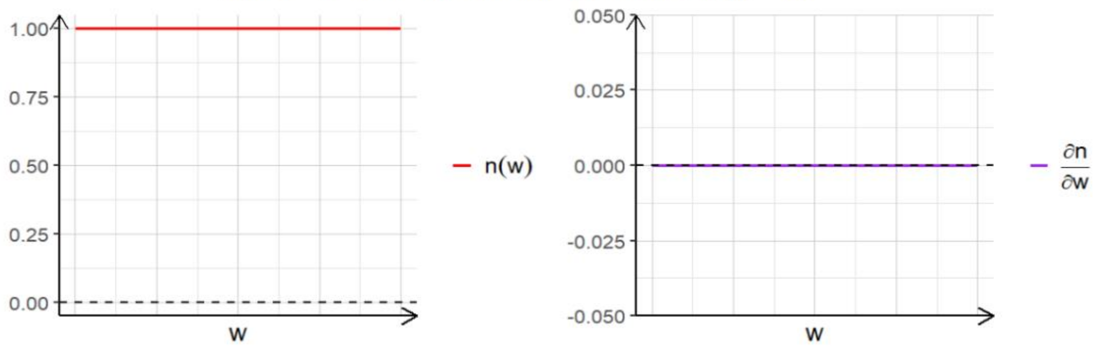

Figure 9: *Plotted Simulation of n_t over w ; ($N_{A,t} \in [1, N-1]$)*

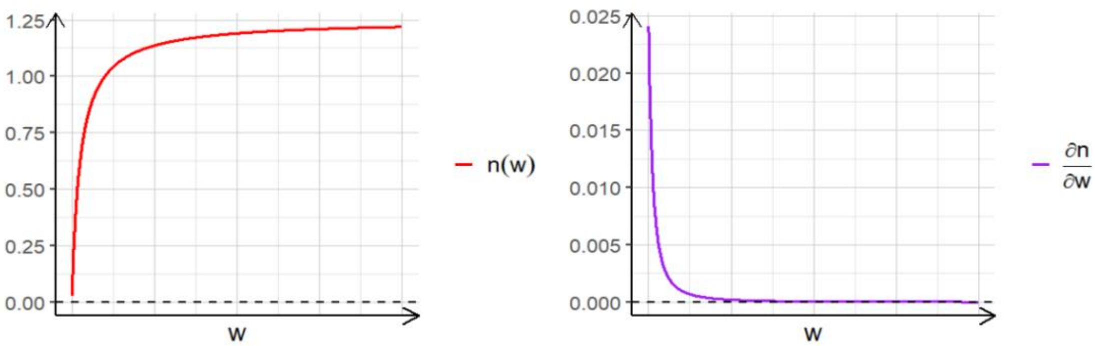

Figure 11: *Plotted Simulation of n_t over w ; ($N_{A,t} = N$)*

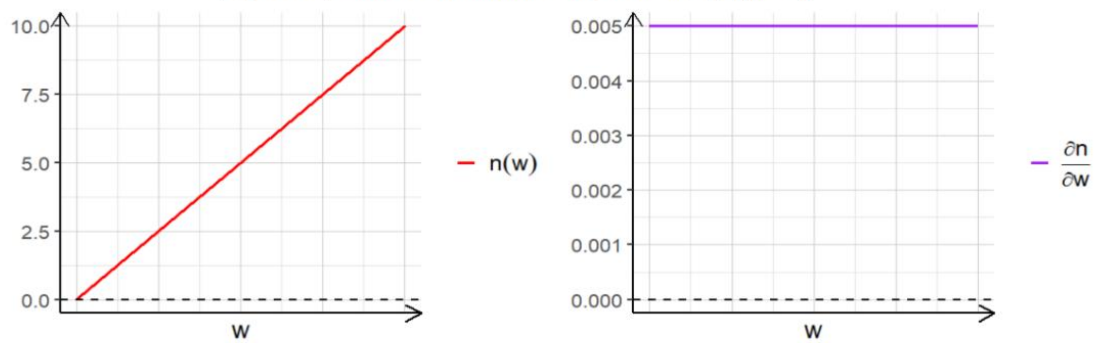

4.2.5. *Automation* ($N_{A,t}$)

$$\frac{\partial n_t}{\partial N_{A,t}} = \frac{\delta}{1+\delta} \cdot \frac{w(-w\phi + \psi)}{\left[w\phi\left(N - N_{A,t}\right) + \psi N_{A,t}\right]^2} \tag{32}$$

Equation (32) is a pivotal takeaway from our model's setup because it implies that the impact of automation on fertility relies on the cost of automating tasks ψ relative to the wage adjusted time-cost $w\phi$ of those same tasks if they are performed manually:

$$\frac{\partial n_t}{\partial N_{A,t}} \text{ is } \begin{cases} > 0 \text{ if: } & w\phi < \psi \\ = 0 \text{ if: } & w\phi = \psi \\ < 0 \text{ if: } & w\phi > \psi \end{cases} \tag{33}$$

Essentially, if the wage-adjusted time cost $w\phi$ exceeds the monetary cost ψ of automation, expanding the automation boundary $N_{A,t}$ reduces the overall cost of childrearing and increases fertility (Figure 11). Conversely, if automation is relatively expensive ($w\phi < \psi$), then further automation may lower fertility by raising household childcare costs (Figure 13). If both costs are equal, then automation has no impact on fertility at all (Figure 12).

Figure 11: *Plotted Simulation of* n_t *over* $N_{A,t}$; ($w\phi > \psi$)

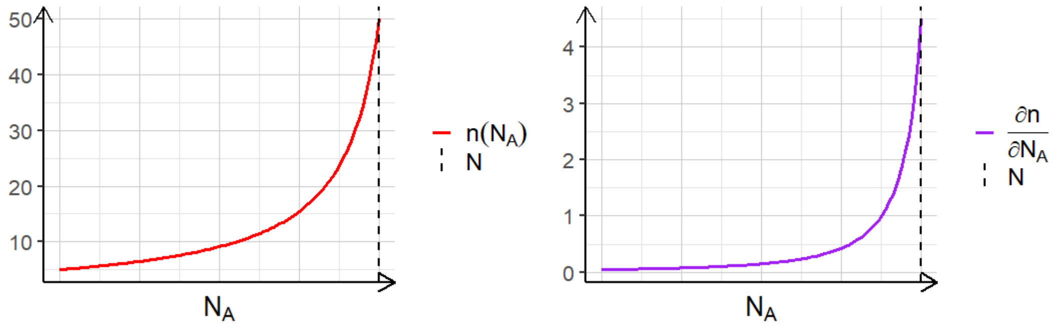

Figure 12: *Plotted Simulation of* n_t *over* $N_{A,t}$; ($w\phi = \psi$)

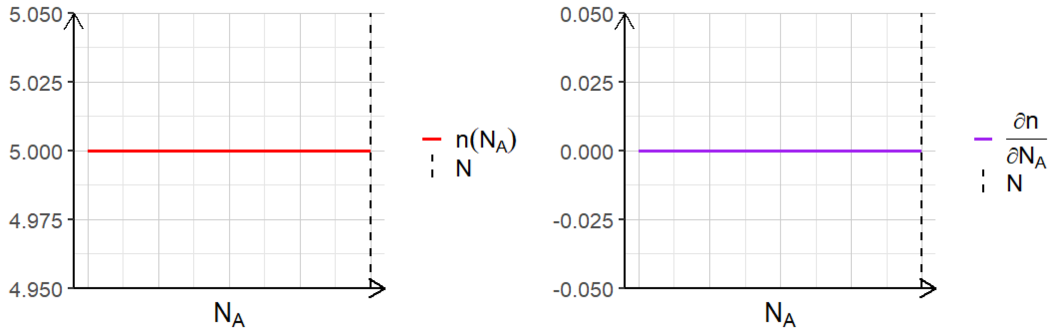

Figure 13: *Plotted Simulation of* n_t *over* $N_{A,t}$; ($w\phi < \psi$)

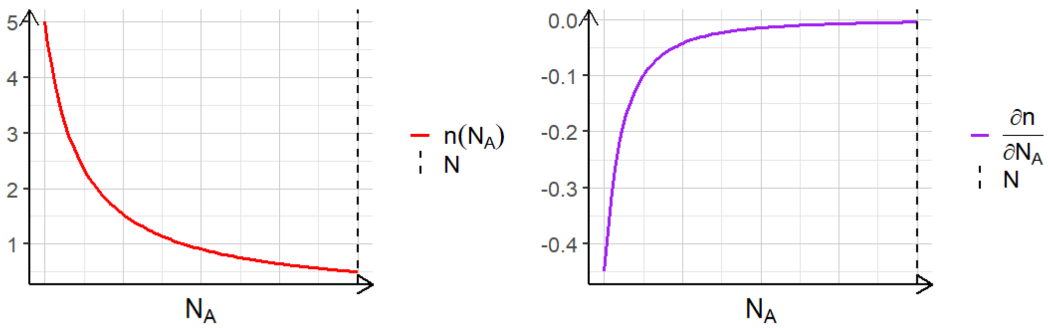

5. Model Discussion

5.1. Engaging with Deopke et al. (2023)

In Doepke et al. (2023),[9] the fall in fertility is partly explained as resulting from an increase in female wages. This is achieved by setting up the model with heterogeneous female and male wages $w_f < w_m$. Since women earn a lower wage, they are assumed to perform all childrearing tasks and use whatever time left to work, while men work full time. This results in the following solution to the household's utility maximisation problem:

$$n = \frac{\delta}{1+\delta} \cdot \frac{1}{\phi} \cdot \left[1 + \frac{w_m}{w_f}\right] \tag{34}$$

where we can see intuitively that closing the gender pay gap reduces fertility while widening the gap raises it. This conclusion, however is based on the questionable assumption that if the household's wife earns less than her husband – even marginally – she will perform **all** the household's childrearing tasks. We recommend future research to explore a different, possibly more nuanced setup in which childrearing tasks are exogenously split between husband and wife proportionately to the female-male wage ratio.

Nonetheless, setting up our model with this assumption is both simple and pertinent,[10] because it allows us to evaluate the conditions needed for automation to mitigate the fall of fertility rates. Doing so results in a fertility solution that somewhat resembles Section 4.2. of the Deopke et al. paper on the "Marketization of Childcare":

$$n_t = \frac{\delta}{1+\delta} \cdot \frac{w_m + w_f}{N_{A,t}\psi + (N - N_{A,t})w_f\phi} \tag{35}$$

Raising male wages here unambiguously raises fertility while raising the female wage has ambiguous effects. We take the partial differential with respect to w_f and set the FOC:

$$\frac{\partial n_t}{\partial w_f} = \frac{\delta}{1+\delta} \cdot \frac{N_{A,t}\psi - (N - N_{A,t})w_m\phi}{\left[N_{A,t}\psi + (N - N_{A,t})w_f\phi\right]^2} = 0 \tag{36}$$

We study how $\frac{\partial n_t}{\partial w_f}$ behaves in function of $N_{A,t}$ (Appendix A.3.) we note that under the conditions $\psi < w_f\phi$ and $w_f \leq w_m$:

$$\frac{\partial^2 n_t}{\partial w_f \partial N_{A,t}} > 0, \quad \lim_{N_{A,t}\to 0} \frac{\partial n_t}{\partial w_f} = \frac{-\delta w_m}{(1+\delta)N w_f^2 \phi} < 0, \quad \text{and} \quad \lim_{N_{A,t}\to N} \frac{\partial n_t}{\partial w_f} = \frac{\delta}{(1+\delta)N\psi} > 0 \tag{37}$$

$$\frac{\partial n_t}{\partial w_f} = 0 \quad \text{when} \quad N_{A,t} = N_A^* = \frac{N w_m \phi}{\psi + w_m \phi} \tag{38}$$

[9] Specifically in Section 2.4.

[10] Though not part of the original model explanation as it distracts from its main purpose.

With no automation, the effect of increasing female wages is negative. As automation kicks in, this effect becomes increasingly positive and, as the automation boundary reaches full automation, increasing female wages increases fertility. There then exists a critical level of technology N_A^* where the negative effect of increasing female wages on fertility becomes positive, as illustrated below (Figure 14).

Figure 14: *Plotted R Simulation of* $\frac{\partial n_t}{\partial w_f}$ *over* $N_{A,t} \in [0, N]$

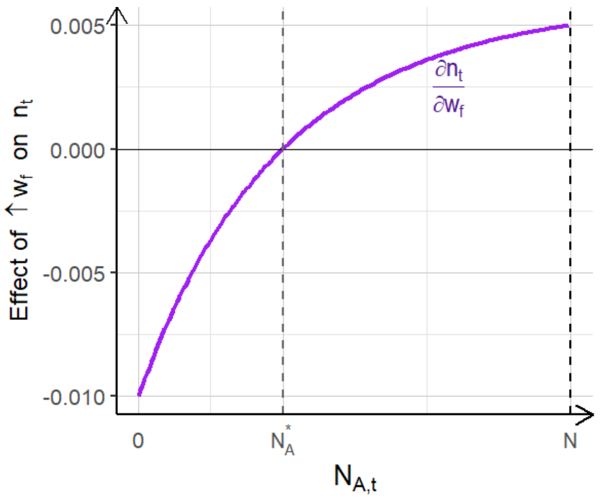

By that logic, in the long-term, whether or not automation can fully reverse the phenomenon described by Doepke et al. (2023), depends on the proportion of childrearing tasks that are possible to automate Δ (Equation 6) satisfying the following condition:

$$\Delta > \frac{N_A^*}{N} \quad \Rightarrow \quad \Delta > \frac{w_m \phi}{\psi + w_m \phi} \tag{40}$$

5.2. Limitations & Future Theory Research

5.2.1. Adoption of Automation

A core simplifying assumption of this model is that households universally and immediately adopt childcare automation as soon as it becomes technologically feasible. In formal terms, the automation boundary $N_{A,t}$ expands smoothly as a function of the technology level A_t (Equation 6), determined by the interaction of technological progress $A_t^{\sigma_y}$, the remaining share of unautomated tasks $1 - N_{A,t}/\Delta N$, and the activation term $1 - e^{\kappa(A_{\text{crit}} - A_t)}$ that governs when automation begins. This formulation fully captures the supply-side potential of automation but abstracts from the demand-side dynamics of household adoption.

In reality, automation adoption is neither universal nor frictionless.[11] Economic, social, and psychological factors strongly shape which households choose to adopt childcare automation and when. High-income households may be early adopters due to affordability; lower-income households may face prohibitive costs. Likewise, parents may resist automating certain tasks they view as emotionally or developmentally critical, regardless of the technology's availability. As a result, the effective automation boundary influencing household fertility decisions may lag well behind the maximum $N_{A,t}$ implied by technological progress alone.

[11] As is argued by Gries and Naudé's (2022) "human services" partial equilibrium model of automation in the labour market.

This distinction matters because the fertility solution (Equation 14) is highly sensitive to the distribution of childcare costs across time and money. If automation reduces the average time cost only partially or unevenly across households, the aggregate fertility response will be smaller than predicted by the current model. Put differently, the model's positive automation–fertility–growth linkage critically depends on widespread and affordable adoption, not just on technological readiness.

Future iterations of this model could endogenize adoption by introducing household heterogeneity and adoption decisions. For example, the model could include an adoption probability p_i, determined by household income, preferences, or cultural attitudes, and define the effective automation boundary $N_{E,t}$ as $N_{E,t} = \sum_{i=1} N_{A,t} p_i$. This would allow the model to capture not just the supply of automation but also its diffusion, possibly yielding more realistic predictions. Additionally, modeling network effects – where the likelihood of adoption increases as more households adopt – would reflect real-world innovation diffusion patterns.

In summary, the current framework likely overstates automation's demographic impact by assuming immediate and uniform adoption. Relaxing this assumption could generate richer, policy-relevant predictions on how affordability, inequality, and social acceptance mediate the relationship between automation, fertility, and growth.

5.2.2. *Improving the link between $\boldsymbol{n_t}$ and $\boldsymbol{L_t}$*

As it stands, our model's link between fertility decisions n_t and total workforce is slightly oversimplified. Specifically, we define n_t as the number of children per household at time t, and we link this directly to population growth (Equation 5) without explicitly modelling the time-lag between childbearing and entry into the workforce.

This effectively treats the model period as spanning roughly one generation (≈ 18 years), which is reasonable considering the scope of this dissertation being focused on the long-run effects of technology on population growth. However, it means our current version of the model may be overlooking shorter-term nuances that could impact fertility. For example, a gradual societal shift where children take increasingly long to become financially independent would increase the time duration of childrearing, and therefore its burden, potentially dampening fertility decisions.

For a more accurate representation, future research could extend our framework to improve the link between the household utility maximisation problem and workforce. One suggestion we propose as an example would be to model household optimization over childbearing rather than total children per household. Households would have the capacity of creating one child per period t, which they will have to care for over an exogenously specified duration τ before that child transitions into the workforce. We could then express fertility as the probability b_t that a household would choose to produce a child during period t. In that case, population growth becomes:

$$\frac{\dot{L_t}}{Lt} = b_{t-\tau} - m \quad ; \quad n_t = \sum_{i=t-\tau}^{t} b_i \tag{41}$$

Households still derive utility from their number of children and consumption. However, the household's budget constraint would need to be re-designed to express the costs b_t as spanning over τ periods rather than just one. Such an approach could allow richer exploration of short-term population and labour force dynamics.

6. Conclusion & Policy Implications

This dissertation has developed a novel theoretical framework to examine how the automation of childcare may affect the fertility decisions that drive long-run economic growth. By re-designing Zeira's task-based approach and Acemoglu and Restrepo's automation frontier to fit the context of parenting, it creates a new link between technology production and the household dynamics of fertility. It then connects household fertility decisions back into SEG population growth dynamics, establishing a feedback loop in which technology – under certain conditions – can sustain the population growth needed to maintain long-run innovation, and economic growth along a balanced growth path.

This research offers two key contributions:

(1) A novel mechanism through which technological progress can *directly* influence population growth by automating time-intensive parenting tasks and converting them into monetary expenditures, thereby altering fertility decisions and potentially establishing supporting levels of demographic expansion consistent with long-run economic growth.

(2) A policy-relevant insight: If we believe that, as explained in Doepke et al (2023)., rising females wages is at the root of the global fertility decline threatening economic growth, the automation of childcare could be the solution to the "empty planet" problem. By potentially reversing the negative effect of rising female wages on fertility to a positive one, automation turns an economic concern into a source of economic growth, while simultaneously promoting gender equality.

The main policy implication of (2) being that lawmakers should move beyond traditional family policies such as tax credits or parental leave, and recognize the role that household technology can play in supporting fertility. Encouraging innovation in affordable childcare automation, reducing regulatory barriers, and disparities in access are key policy measures that could enhance the demographic and economic benefits of technological progress.

However, the deliberately theoretical nature of this paper implies that demonstrating that childcare automation would reverse current fertility trends hinges on empirically validating the plausibility of conditions required for the model's predictions to hold. Until such evidence is available, the policy implications outlined above should be viewed as exploratory rather than actionable.

In closing, this dissertation highlights an underexplored but potentially transformative link between automation, fertility, and growth. As advanced economies grapple with the twin challenges of demographic decline and technological disruption, understanding how these forces interact will be central to designing policies that promote both economic dynamism and social resilience.

Appendices

A.1. Summary of Model Setup

Growth Dynamics

Final Good Production:

$$Y_t = A_t^{\sigma_y} L_{Y,t}$$

Idea Production:

$$\frac{\dot{A}_t}{A_t} = z L_{A,t} A_t^{-\beta}$$

Population Dynamics:

$$\frac{L_{A,t}}{L_{Y,t}} = l \qquad L_t = L_{A,t} + L_{Y,t}$$

Population Growth:

$$\frac{\dot{L}_t}{L_t} = \frac{\dot{n}_t}{2n_t} - m$$

Automation Boundary

$$\dot{N}_{A,t} = A_t^{\sigma_n}\left(1 - \frac{N_{A,t}}{\Delta \mathrm{N}}\right)\left(1 - e^{\kappa(A_{\mathrm{crit}} - A_t)}\right)$$

Household Childbearing Decisions

Utility Function:

$$u(c, n) = \log(c) + \delta \log(n)$$

Time and Budget Constraints:

$$\lambda_t = 1 - n_t \phi\left(N - N_{A,t}\right) \;\; ; \;\; w\lambda_t = c_t + n_t \psi N_{A,t}$$

Maximisation Problem:

$$\max_{c_t, n_t}: \log(c_t) + \delta \log(n_t) \quad s.t. \quad w\left(1 - n_t\left(N - N_{A,t}\right)\phi\right) = c_t + n_t N_{A,t} \psi$$

Closed-Form Fertility Solution:

$$n_t = \frac{\delta}{1 + \delta} \cdot \frac{w}{w\phi\left(N - N_{A,t}\right) + \psi N_{A,t}}$$

A.2. Fertility Solution Derivation

We start from the utility maximisation problem:

$$\max_{c_t, n_t}: \log(c_t) + \delta \log(n_t) \quad s.t. \quad w\big(1 - n_t(N - N_{A,t})\phi\big) = c_t + n_t N_{A,t}\psi$$

Isolate c_t:

$$w\big(1 - n_t(N - N_{A,t})\phi\big) = c_t + n_t N_{A,t}\psi \quad \Leftrightarrow \quad c_t = w\big(1 - n_t(N - N_{A,t})\phi\big) - n_t N_{A,t}\psi$$

Substitute into the objective function:

$$\max_{n_t}: \log\big(w\big(1 - n_t(N - N_{A,t})\phi\big) - n_t N_{A,t}\psi\big) + \delta \log(n_t)$$

$$\text{Define:} \quad C(n_t) = w\big(1 - n_t(N - N_{A,t})\phi\big) - n_t N_{A,t}\psi$$

Differentiate with respect to n_t and set the F.O.C.

$$\frac{\partial}{\partial n_t}\big[\log\big(C(n_t)\big) + \delta \log(n_t)\big] = \frac{-w(N - N_{A,t})\phi - N_{A,t}\psi}{C(n_t)} - \frac{\delta}{n_t} = 0$$

Solving for n_t:

$$\frac{-w(N - N_{A,t})\phi - N_{A,t}\psi}{C(n_t)} = -\frac{\delta}{n_t}$$

$$\Rightarrow \quad \big(-w(N - N_{A,t})\phi - N_{A,t}\psi\big)n_t = -\delta C(n_t)$$

$$\Rightarrow \quad \big(-w(N - N_{A,t})\phi - N_{A,t}\psi\big)n_t = -\delta\big[w\big(1 - n_t(N - N_{A,t})\phi\big) - n_t N_{A,t}\psi\big]$$

Expand and group by n_t terms:

$$\Rightarrow \quad \big[-w(N - N_{A,t})\phi - N_{A,t}\psi - \delta w(N - N_{A,t})\phi - \delta N_{A,t}\psi\big]\, n_t = -\delta w$$

Isolate n_t:

$$n_t = \frac{\delta w}{(1+\delta)\big[w\phi(N - N_{A,t}) + \psi N_{A,t}\big]}$$

Rearrange for the final solution:

$$\boxed{n_t = \frac{\delta}{1+\delta} \cdot \frac{w}{w\phi(N - N_{A,t}) + \psi N_{A,t}}}$$

A.3. Engaging With Doepke et al. (2023) Derivation

Recall:

$$n_t = \frac{\delta}{1+\delta} \cdot \frac{w_m + w_f}{N_{A,t}\psi + (N - N_{A,t})w_f\phi}$$

and

$$\frac{\partial n_t}{\partial w_f} = \frac{\delta}{1+\delta} \cdot \frac{N_{A,t}\psi - (N - N_{A,t})w_m\phi}{\left[N_{A,t}\psi + (N - N_{A,t})w_f\phi\right]^2} = 0$$

We take the cross-partial derivative of n_t with respect to w_f and $N_{A,t}$:

$$\frac{\partial^2 n_t}{\partial w_f \partial N_{A,t}} = \frac{\delta}{1+\delta} \cdot \frac{(\psi + w_m\phi)\left[N_{A,t}\psi + (N - N_{A,t})w_f\phi\right]^2 - \left[N_{A,t}\psi - (N - N_{A,t})w_m\phi\right] \cdot 2\left[N_{A,t}\psi + (N - N_{A,t})w_f\phi\right](\psi - w_f\phi)}{\left[N_{A,t}\psi + (N - N_{A,t})w_f\phi\right]^4} > 0$$

$$\Rightarrow (\psi + w_m\phi)\left[N_{A,t}\psi + (N - N_{A,t})w_f\phi\right]^2 - \left[N_{A,t}\psi - (N - N_{A,t})w_m\phi\right] \cdot 2\left[N_{A,t}\psi + (N - N_{A,t})w_f\phi\right](\psi - w_f\phi)$$

$$\Rightarrow (\psi + w_m\phi)\left[N_{A,t}\psi + (N - N_{A,t})w_f\phi\right] > 2\left[N_{A,t}\psi - (N - N_{A,t})w_m\phi\right](\psi - w_f\phi)$$

Which holds under the conditions:

$$\psi > w_f\phi \quad \text{and} \quad w_m \geq w_f$$

In that case, the following is assured:

$$\boxed{\frac{\partial^2 n_t}{\partial w_f \partial N_{A,t}} > 0}$$

A.4. Simulation Scripts

All scripts can be accessed by clicking the GitHub button to access the script online.

Evolution of the Automation Boundary $N_{A,t}$

Figure 2

The `AutoBoundary_TechLvl.py` simulation script was coded on Python when I was designing the automation boundary function (around Jan 2025), as a way to test that it behaved as expected.

Fertility Solution (n_t) Comparative Statics

All coded during the writing process of the dissertation (around April 2025) to serve as illustrations of the comparative statics analysis of fertility solution equation (14). I switched from Python to R code for higher code efficiency.

There is one script per parameter/variable studied. Each script defines n_t as a function of the studied parameter/variable and – ceteris paribus – computes the values for both n_t and its partial differential with respect to the studied parameter/variable. Different outputs are obtained by changing the parameter settings at the start of each script.

- Childrearing Preference δ: `sim_n(delta).R`
  Figure 3
- Avg. Time Cost of Tasks ϕ: `sim_n(phi).R`
  Figures 4 and 5
- Avg. Monetary Cost of Tasks ψ: `sim_n(psi).R`
  Figures 6 and 7
- Wages w: `sim_n(w).R`
  Figures 8, 9, and 10
- Automation $N_{A,t}$: `sim_n(N_At).R`
  Figures 11, 12, and 13

Engaging with Doepke et al. (2023)

Figure 14

The `sim_partialDiff(N_At).R` script simulates the behaviour of the partial derivative of n_t with respect to w_f as a function of $N_{A,t}$

Resources

Bibliography

Acemoglu, D. and Restrepo, P. (2018). The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment. *American Economic Review*, 108(6), pp.1488–1542.

Acemoglu, D. and Restrepo, P. (2020). Robots and Jobs: Evidence from US Labor Markets. *Journal of Political Economy*, [online] 128(6), p.2188-2244. doi:https://doi.org/10.1086/705716.

Acemoglu, D. and Restrepo, P. (2022). Tasks, Automation, and the Rise in U.S. Wage Inequality. *Econometrica*, 90(5), pp.1973–2016. doi:https://doi.org/10.3982/ecta19815.

Aghion, P., Antonin, C. and Bunel, S. (2019). Artificial Intelligence, Growth and Employment: The Role of Policy. *Economie et Statistique / Economics and Statistics*, (510-511-512), pp.149–164. doi:https://doi.org/10.24187/ecostat.2019.510t.1994.

Broadbent, E., Stafford, R. and MacDonald, B. (2009). Acceptance of Healthcare Robots for the Older Population: Review and Future Directions. *International Journal of Social Robotics*, 1(4), pp.319–330. doi:https://doi.org/10.1007/s12369-009-0030-6.

Doepke, M., Hannusch, A., Kindermann, F. and Tertilt, M. (2023). The Economics of fertility: a New Era. *Elsevier eBooks*, 1(1), pp.151–254. doi:https://doi.org/10.1016/bs.hefam.2023.01.003.

Gries, T. and Naudé, W. (2022). Modelling artificial intelligence in economics. *Journal for Labour Market Research*, 56(1). doi:https://doi.org/10.1186/s12651-022-00319-2.

Jones, C.I. (1995). R & D-Based Models of Economic Growth. *Journal of Political Economy*, 103(4), pp.759–784. doi:https://doi.org/10.1086/262002.

Jones, C.I. (2020). The End of Economic Growth? Unintended Consequences of a Declining Population. *American Economic Review*, 112(11), pp.3489–3527. doi:https://doi.org/10.1257/aer.20201605.

Jones, C.I. (2022). The Past and Future of Economic Growth: A Semi-Endogenous Perspective. *Annual Review of Economics*, 14(1), pp.125–152. doi:https://doi.org/10.1146/annurev-economics-080521-012458.

Luo, Q.Z. and Hsiao-Chin, L.Y. (2023). The Influence of AI-Powered Adaptive Learning Platforms on Student Performance in Chinese Classrooms. *Journal of Education*, 6(3), pp.1–12. doi:https://doi.org/10.53819/81018102t4181.

Parents. (2024). *Why Gen Z Is Calling Today's Kids 'iPad Kids'*. [online] Available at: https://www.parents.com/what-are-ipad-kids-8692488.

United Nations (2024). *World Population Prospects 2024: Summary of Results*. [online] DESA Publications. Available at: https://desapublications.un.org/publications/world-population-prospects-2024-summary-results?_gl=1.

Wei, W. and Xie, D. (2022). Can AI Help to Raise the Fertility Rate? *SSRN Electronic Journal*. doi:https://doi.org/10.2139/ssrn.4171328.

Zeira, J. (1998). Workers, Machines, and Economic Growth. *The Quarterly Journal of Economics*, 113(4), pp.1091–1117. doi:https://doi.org/10.1162/003355398555847.

Use of AI Content Declaration

The Student (u216980) acknowledges and declares that during the preparation of this assignment, Artificial Intelligence was used exclusively for the following:

- **Help with code syntax and simulation graphics formatting** (Model: ChatGPT 4o)
- **Minor rewording suggestions to improve flow and academic tone** (Model: ChatGPT-"Academic Writing"; Author: Moahaimen Talib)

No content was generated or inserted directly into the final submission without critical review and substantial modification.
The student confirms that all research, knowledge, ideas, analysis, and structure presented in this assignment are entirely his own, and that this submission fully aligns with the objectives of the assessment and the University's academic integrity standards.