Ca' Foscari
University
of Venice

Master Degree

in

Economics, Finance and Sustainability

Curriculum: Economics QEM

Final Thesis

When Automation affects Labour Productivity

A growth model of human-machine coexistence

Supervisor
Ch. Prof. Pietro Dino Enrico Dindo

Graduand
Elia Paccagnella
Matriculation Number 879028

Academic Year
2024 / 2025

Abstract

This thesis presents a task-based sustained growth model where human labour and physical capital compete as production factors. The model introduces to the discussion on the potential positive effect that automation has on labour productivity.

To begin, chapter one will offer an introduction on automation with the analysis of the historical effects of automation on employment.

The second chapter will be dedicated to the definition of the economic model. Section one is dedicated to a review of the task-based approach to the economics of automation. This section will provide a generalization of task-based based models as well as allowing to choose the best model to start with. Section two presents the main contribution of this thesis: the definition a task-based automation-driven growth model.

The third chapter concludes with an historical and technical focus on Artificial Intelligence with a final consideration of the potential effects of an Artificial General Intelligence on economic modelling.

Contents

Chapter 1

An introduction to human-machine interaction

We are now experiencing a rapid development of Artificial Intelligence (AI) with application in many, previously unexpected, sectors. ChatGPT - the know-it-all AI chatbot developed by OpenAI - became the fastest growing consumer application in the world in terms of active users at its launch[1]. History shows that, in moments of fast technological advancements like these, humans tend to start doubting their abilities to compete with machines in production. This is similar to what many think also now seeing the effects of AI in the labour market. The rapid pace at which this technology is - at least apparently - advancing, raised huge concerns, both in the workforce and in the media, regarding the imminent threat of technological job substitution. This thesis aims at investigating this phenomenon from an economic point of view.

Automation has been a core component of human labour since the development of the very first agricultural techniques. Mankind did not want to irrigate constantly crops and invented canals. Equally, grounding wheat by hand was tedious and farmers invented mills. Automating tedious tasks is a natural human desire.

[1] Hu, 2023.

Even so, automation began to raise concerns in the workforce with the First Industrial Revolution. During that period, not only technological advances allowed for task automation, but also they enabled a completely new organization of the workforce. Many rural workers moved to the cities looking for higher wages. Such high wages could be interpreted as a direct effect of automation as we will analyse in this chapter. Moreover, automation allowed to rump up production and leverage economies of scale. And with Henry Ford's Model T, factory workers became also the primary customer segment for volume of sales opening up new target markets. It could be inferred, thus, that automation had a significant indirect effect also on the demand side. Is it possible to define a globally-applicable and timeless economic model describing the dynamics of the labour market caused by automation?

Many would to think at automation as an evil entity looking for workers to lay off. The current perception of AI delivered by many newspapers' headlines is not so different. However, literature proves this definition to be wrong at least twice. Firstly, it is now common practice in economic literature to consider the effects automation has at a task level and not at a job level thanks to the contribution of Acemoglu and Autor, 2011. We define a job as the process needed to deliver a final output from a raw input. Any job can be split in multiple tasks representing every activity needed to complete a job. Also tasks can be seen as a function taking some raw and/or pre-processed inputs and delivering a semifinished output. This distinction is crucial since it allows us to develop an economic model that closely replicates what we observe from the real world data. Citing Bowen, 1966 "technology eliminates jobs and not works". Following our definitions Bowen's sentence can be translated as: automation replaces workers in carrying out specific tasks but not in doing their job as a hole.

Secondly, as we empirically discussed in the introduction, automation can either have positive or negative effects on the labour share for a specific task. Following this observation, Autor, 2015 highlights three aspects that should allow us to define the effects that technology-driven automation has historically had on the workforce.

1. **Substitutability of automated tasks** — A given task can either substitute or complement for another task. In the case of complementarity, when one task is automated - thus made more efficient - we expect an increase in the demand also for the other task. When, instead, two tasks are substitutes, if one gets automated, the demand for the other one will drop.

2. **Elasticity of labour supply** — The effects of task complementarity on wages can either be exacerbated or dumped given the elasticity of labour supply. When elasticity is close to zero, the automation of a task is expected to generate huge wage gains for workers carrying out complementary tasks.

3. **Output/Income elasticities of demand** — Ultimately demand increases are the ones that allow for wage gains. Provided that automation generates an increase in output, such increase affects demand in different ways based on the output elasticity of demand. The same will hold for effects on income.

An example presented by Autor, 2015 is the effect Automated Teller Machines (ATMs) had on the demand for banking services. Contrary to the contemporary beliefs, ATMs increased wages of and demand for bank tellers in the US. This is due to the fact that ATMs automated only one task carried out by bank tellers: distributing cash. This type of automation allowed bank teller to dedicate more time to complementary tasks, creating a more personalized experience for their customers. In the meanwhile, it lowered the overall costs opening up banking services opening to less wealthy customers. This increased demand for banking services in turn generated wage gains for bank tellers since their job was more requested. The elasticity of job supply was low provided that only few people had the education necessary to perform well as bank teller.
A counter example in the same sector might be the automation enabled by the digitalization of banking services. In the most recent years, we are observing the disruption of the traditional banking sector due to electronic banks (*e-banks*). Digitalization allowed e-banks to automate and/or centralize traditional banking services up to the point that a physical location

is no longer needed for many customers. An European example is the English bank Revolut born in 2015 which forced many other traditional banks to open online-only branches. It is quite early to find empirical studies investigating the effects of e-banks on banks' employment. Some studies on the topic suggest digitalization of the banking sector can be correlated with a lower concentration of the market[2]. Looking at a similar disruption in the market for phone carrier services, however, would prove that complete virtualization of the service is not possible until technology will be able to automate also tasks requiring deep human interaction and empathy.

These examples highlight the intricacy of the effects automation can have on employment. However, they could also be used to observe some common characteristics defining which tasks can or cannot be automated. On this note, let us consider automatability of a given task i to be a probability distribution $a_i(z \mid A, x)$ over levels of automatability z with parameter A representing the current technology level. In this formalization, we allow the sample space Ω to be either binary ($\Omega \equiv \{0, 1\}$) - so that z can represent only automatable and non-automatable tasks - or continuous ($\Omega \equiv [0, 1]$). We define also that, keeping z constant, a_i must be weakly increasing in A. Looking at the past, our $a_i(z \mid A, x)$ is no longer a probability distribution function but rather it represents its actualization in the labour market. For this reason, looking at historical data we can focus on the binary sample space and determine the variables that define the distinction between automatable and non-automatable task. This is the empirical approach that Autor, 2015 takes to define automatability in the Information Technologies (IT) era. Computers allowed for the automation of tasks through the definition of a list of precise and specific instructions written in computer code. This means that highly repetitive tasks carried out in a controlled and predictable environment suddenly became highly automatable. On the other hand, computer scientists quickly realized their limitations in describing more complex, learning-based actions through code. For this reason, managerial tasks were hardly automatable.

[2] Ghasemi et al., 2021.

Duckworth et al., 2019 present an advancement in the definition of automatability providing an analytical framework. Through the categorization of working activities in three levels (work activities w, tasks x, occupations o), they develop a robust method to estimate the (perceived) automatability of a task starting from experts' opinions. Their finding are similar to the ones of Autor in terms of distribution of task automatability. One could argue that their scientific and analytical approach could be biased due to its own nature. If automatability is evaluated only based on experts' opinions, it could be that exposure of the expert group to the latest advancement in technology shifts their estimations of automatability higher than the actual value. Nevertheless, based on the *wisdom of the crowd* principle, we can still believe this method to generate a good proxy for automatability.

Chui et al., 2015 and then Chui et al., 2016 also investigate automatability potential at a task level for numerous industries. Their methodology, however, is not clearly explained but reinforces the results of the previous papers. However, their study goes further in defining which technological innovations will have the highest impact on task automatability. They state that the development of machine natural language understanding could vastly increase automatability in retail, finance and insurance, and healthcare. Ten years after their publication, one could state that this technological innovation has both been developed and brought to market but still needs to be implemented in the automation of many of those tasks. This observation captures the difference that exists between task automatability and their automation. Actual automation of a task takes place only if automation gains exceed its costs. This different will be central in economic modelling.

Unfortunately, it is impossible to provide a unique definition since automatability is a function of technology available at a specific time period. This intuition is essential in understanding the effects of automation on the labour market. Autor, 2015 hints something when discussing the characteristics a task should have to be automated though the use of Information and Communication Technologies (ICT). Computers are used to interpret computer code and convert it into actions. The main limitation, in this case, is thus the definition of computer code that allows for a certain action (task) to be automated. While this could not look like a huge limitation in the first place, in reality it is. ICT engineers can only write code for actions that they are able to describe with a sequence of instructions. While user inputs allow for more flexibility also in uncontrolled environments, there exist some actions that cannot be converted into computer code. In this case, ICT define automatable tasks as repetitive, predictable, carried out in a controlled environment. On the other hand, non-automatable tasks are emotional, empathy-based, or carried out in an unpredictable environment. Every technological advancement could, thus, change the definition of automatability. A reasonable assumption could be that such a change is monotonically quasi-increasing, in the sense that technological advancements can either increase or leave unchanged the size of the set of automatable tasks.

It is from a clearly-bounded definition of automatability that we can understand how technology impacts employment. We can still follow Autor, 2015 in examining the case of ICT automation. Following the definition of automatability under ICT and ranking tasks in three groups respectively requiring low, medium, and high expertise, we could state that ICT has an high potential of automating medium-skilled tasks while leaving unaffected low- and high-skilled tasks. This is referred to as the *theory of employment polarization*: advancements in technology or further integration of existing technologies in the workforce are expected to lower the demand for medium-skilled tasks. Medium-skilled workers will thus be displaced in low-skilled, high-skilled, or newly born positions. While the shift from a medium-skilled position to a lower-skilled one can be virtually instant, moving a medium-skilled worker to a higher-skilled position requires a retraining period. As for what regards new positions, their effect is highly dependent on the assumptions we make on their birth and development. Autor, 2015 shows *employment polarization* can be observed between 1979 and 2012 in US employment data. Most importantly, they prove that *employment polarization* does not directly lead to *wage polarization*. It must be noted, however, that the polarization effect is extremely reliant on the assumption defined by the automation technology taken into consideration - IT - works. When the automation technology considered changes, the polarization effect is not proved to hold strong.

Let me imagine, only for a moment, that there exists a technology so advanced that can virtually automate any task in a cost-effective way. In such a world, human labour would no longer be required. Food is produced and delivered automatically to your home, scientific advancements autonomously flourish, and even politics are automated. Then, citing Autor, 2015, “if human labour is indeed rendered superfluous by automation, then our chief economic problem will be one of distribution, not of scarcity”. Economists are recently attempting to model also an economy of this kind with some remarkable examples in Restrepo, 2025 and Bara, 2025.

As last remark, we should also consider which aspects - if any - rule technology development. To summarize the literature advancements on this regard, I suggest to bound technology advancements to four main constraints:

1. an economic constraint: innovations are developed only if the costs in research and development required can be feasibly recovered either through competitive advantage or the creation of an entirely new market;

2. an environmental constraint: production (even of services) requires access to natural resources. These resources are, many times, exhaustible and not renewable. If a technology consumes too many resources, the negative spillover effects it generates begin to hold the development of such technology back;

3. a technological constraint: an idea need the underling technology to be implemented. Technology can, thus, hold technological advancement back. Iterative innovations increase the technological threshold only slightly while radical innovations increase it drastically. The technological threshold can also be lowered from government bans/restrictions;

4. a sociocultural constraint: society must accept technology to integrate it into their lives.

The next section will investigate how advancements in Artificial Intelligence could be limited by any of the afore-mentioned constraints.

Chapter 2

Economic modelling

2.1 A task-based approach to the Economics of Automation

In Section 3.1.2 we argued that Artificial Intelligence is an umbrella term that encompasses any technology that attempts to replicate or simulate behaviour that is considered to be typically human. However, this general definition often is too strict for economic modelling whose goal is to abstract the essence of nature into tractable mathematical relationships between variables. For this reason economists usually abstract the concept of AI even more considering exclusively the effects their interested in for their study. Instead of using narrower technical terms like Convolutional Neural Networks or Generalized Pre-trained Transformers, economist ofter revert to general terms like innovation, automation, robots, or machines. This approach not only allows economists to focus their research on the specific effect caused by AI in the market of interest, but it also generalizes their findings to other technologies that generate similar effects. Every term, however, has a slightly different interpretation of the problem in hand. Speaking about innovation, for example, shifts the focus on the growth potential of the new technology and usually requires entrepreneurial investment to be developed. Automation, instead, is often used when observing the effect of technological improvements on the labour market, including distributional effects. Lastly, robots and machines are mostly used as synonyms for automation even thought they assume that technological automation occurs only for tasks requiring some interaction with the physical world. This section, provides a literature review of the main advancements in the economics of automation mainly focusing on the latest research trend of task-based models.

As highlighted also in **??**, automation is not a newly discovered topic and *technological unemployment* has occasionally been a concern for economists also during the 19th and 20th centuries. Karl Marx was among the firsts to advocate the idea that automation will eventually lead to class disparities. However, some hints of *technological pessimism* can be found also on the previous work of David Ricardo. These are only two examples of some the most influential minds in the field of economics that raised (different) concerns about the effects that automation could have on the labour market and on the global and national wealth distributions. It must be acknowledged, however, that over time counterparts to these views have not been scarce.

The traditional approach to automation in economic modeling has been to consider the effect of factor-specific technologies. By using either a capital-augmenting technology or labour-augmenting one, economists could model the main displacement effect of automation. However, this approach has one main underling assumption: either technological advancements are able to fully replace workers in their jobs, or automation does not take place. Moreover, these models are not able to explain the effects observed in Autor, 2015. A capital-augmenting technology is able to explain workers' displacement from one sector to another but it cannot prove why some works are still performed by humans even thought the automation technology for those works exists and it is economically efficient to implement it.

For this reason, at the beginning of the 21st century, economists began to explore a new approach to automation models. Empirical observation made clear that automation often does not replace workers in carrying out entire jobs but rather single tasks of a particular job. Based on the work of Zeira, 1998 and Costinot and Vogel, 2009, Acemoglu and Autor, 2011 introduced a task-based approach to the economics of automation. Such approach considers the production of a unique final good that requires the contribution of multiple tasks i to be produced. Following the suggestion of Lu and Zhou, 2021 the production function of virtually any model following the task-based approach can be summarized in the following generalized form:

$$Y(t) = G\left(N(t), y_i(t)\right) \tag{2.1}$$

where $N(t)$ represents innovation at time t and $y_i(t)$ is the production function of task $i \in [i_{\text{min}}, i_{\text{max}}]$. A task can be interpreted as an intermediate good that can be produced using a combination of labour and capital following a task-specific production function

$$y_i(t) = A(t) \cdot F\left(\alpha(t)K_i(t), \beta(t)L_i(t)\right) \tag{2.2}$$

Every task-based model must have an assumption on how tasks are ordered in the task space $[i_{\text{min}}, i_{\text{max}}]$. The most common assumption is that tasks are ordered from the easiest to the hardest to automate given current technologies.

Moreover, in all these models there is a common assumption that technology allows only certain tasks to be automated. There exists a set of *automatable tasks* $I \subseteq [i_{\text{min}}, i_{\text{max}}]$. The size of this set determines the automatability potential of the entire economy. To keep the assumptions as generalized as possible, this set I could also be open and non-convex. A closed, convex set however helps in making the model tractable and easier to understand and, many times, non-convexity in I can be solved simply by changing the order of tasks. Within the set I, market forces engage to determine how much labour and capital to use in task production. These forces will determine the set $\tilde{I} \subseteq I$ of *automated tasks*, for which a capital-intensive production process is adopted in equilibrium rather then the labour-intensive alternative.

The consumer side of the economy changes based on the objectives of the model but usually consists of a representative household with a constant relative risk aversion utility function.

Following this generalized model, I will now attempt to summarizes the most influential task-based models.

2.1.1 Zeira (1988) model

This subsection presents the full version of the model introduced in Section III. of Zeira, 1998: an overlapping generations (OLG) model in continuous time.
The production function of the unique final good is

$$\ln Y_t = \ln a + \int_0^1 \ln x_{i,t} di$$

where $A_t = a$ is the total factor productivity and the generic production function $F(\alpha_t K_{i,t}, \beta_t L_{i,t})$ becomes a simple sum of all the tasks' inputs. Note that the author uses this Cobb-Douglas production function for the sake of exposition but results will hold also with other forms. Moreover, observe that the underlying assumption presented in the definition of the production function is that there exists an (infinite) set of tasks indexed $i \in [0, 1]$.
Automatable tasks are defined by a technological threshold f_t such that all tasks $i \in I \equiv [0, f_t]$ are automatable. f_t is increasing in t.
Non-automated tasks (or *manual technologies* in the original paper) are produced using a quantity $l_{0,i}$ of manual labour and $k_{0,i}$ of capital.
Automated tasks (or *industrial technologies* in the original model) instead use respectively $l_{1,i}$ and $k_{1,1}$ quantities such that

$$l_{0,i} > l_{1,i} \quad \text{and} \quad k_{0,i} < k_{1,i}, \quad \forall i \in I$$

or in other words, non-automated tasks are labour-intensive to produce while automated tasks are capital-intensive.
Let

$$h(i) = \frac{k_{1,i} - k_{0,i}}{l_{0,i} - l_{1,i}} > 0$$

Assumption(Z1998) 1 *$h(i)$ is non decreasing in i. Tasks are ordered from the most labour intensive to the most capital intensive.*

Note that this function $h(i)$ could be seen as a measure of the technological level of the economy for the production of task i as well as a measure of the complexity of a task.

On the consumer side, there are two overlapping generations of households:

1. *young households* both consume c_1 and supply one unit of labour;
2. *old households* do not work and consume c_2 units per time period.

For simplicity, the author considers no population growth and normalizes the population size to 1.

All markets are perfectly competitive. There is full capital mobility. Final goods can be sold in a global market while intermediate goods and labour are traded on a locally.

There is a unique equilibrium where technology adoption is defined by a variable

$$m_t = \min\{m^*, f_{t-1}\}$$

where $m^* = \max\{m : w(m) \geq Rh(m)\}$ is the threshold determining the equilibrium set of adopted technologies under unconstrained technology while f_{t-1} represents the technological constraint. Thus, in equilibrium, there could be two cases:

- when $m^* > f_{t-1}$ the economy is technologically constraint;
- when instead $m^* \leq f_{t-1}$ further technological advancements should be focused on reducing $h(m)$ instead of introducing new automation.

The author uses the model to explain the observed difference in technology adoption, the subsequent effect on country productivity, and growth divergence between different countries.

Concluding this brief summary of the model presented in Zeira, 1998, it should be noted that the original paper does not consider is explicitly as tasks but rather as intermediate goods. The difference is nearly absent but still present. Considering the is as tasks makes it easier to argue that the production of one task can be carried out also exclusively using labour or capital. This would allow to simplify the model while retaining the same level of fidelity.

2.1.2 Acemoglu-Author (2011) model

Acemoglu and Autor, 2011 build on the Zeira (1988) model introducing a mechanism of comparative advantage between workers of different skill levels and capital as production factors. For this reason, the authors refer to this model as a Riccardian model of the labour market. They argue that "The distinction between skills and tasks becomes relevant, in fact central, when workers of a given skill level can potentially perform a variety of tasks and, moreover, can change the set of tasks that they perform in response to changes in supplies or technology.". Their model is a static Riccardian model with a finite level of skill groups and exogenous technology.

In this economy there is a unique final good produced following the production function

$$\ln Y = \int_0^1 \ln y(i) di$$

where $y(i)$ is the production function for task i. Note that here they simplify the model normalizing the total factor productivity to 1. Moreover, they use a similar assumption to Zeira, 1998 for task indexing: tasks can be represented with an index $i \in [0, 1]$.

There are three groups of workers differentiated by their skill level. Workers with a low skill level are represented by the letter L, middle-skilled workers M, and highly-skilled workers H. This review considers the simpler version of the model with a fixed inelastic labour supply.

The production function of task i is

$$y(i) = A_L \alpha_L(i) l(i) + A_M \alpha_M(i) m(i) + A_H \alpha_H(i) h(i) + A_K \alpha_K(i) k(i)$$

where

- the A terms are the specific factor-augmenting technologies;
- the α represent the task-specific productivity schedule for low, me-

dium, and high productivity workers as well as for capital. This is the term that defines the comparative advantage of one production factor over the others for a specific task. When a factor of production cannot be used in the production of the task (eg. when a group of workers has insufficient skills to accomplish the task or when there is no technology that allow capital to be used in the production of such task) their α is set to 0;

- $l(i)$, $m(i)$, $h(i)$, and $k(i)$ are respectively the quantities of low, medium, and high skill workers and capital allocated to the production of the task.

Assumption(AA2011) 1 *$\alpha_L(i)/\alpha_M(i)$ and $\alpha_M(i)/\alpha_H(i)$ are continuously differentiable and strictly increasing. This implies that, if we consider highly-skilled workers to have a comparative advantage in more complex tasks with respect to the other production factors, task indexes must be ordered from the least to the most complex.*

The factor market clearing conditions are such that

$$\int_0^1 l(i)\,di \leq L, \quad \int_0^1 m(i)\,di \leq M, \quad \text{and} \quad \int_0^1 h(i)\,di \leq H.$$

There is a unique equilibrium defined as an allocation $(I_L,\ I_H,\ P_L,\ P_M,\ P_H,\ w_L,\ w_M,\ w_H)$ where

- the Ps are the price indices of tasks performed by workers of a specific skill group;
- I_L and I_H represent the thresholds such that every task $i < I_L$ is produced using low skill work, every task $i \in [I_L, I, H]$ is produced by medium skill workers, and every task $i > I_H$ is produced with high skill labour;
- w_L, w_M, and w_H are the wages for low, medium, and high skill workers.

The comparative statics of the model show that the introduction of a

high-skill biased technology will directly reduce I_H, thus allocating more tasks to highly skilled workers, and indirectly reduce I_L since medium skill workers displace by this biased technological increase relocate to lower skill tasks. Similarly, also medium-skill biased technological change and low-skill biased technological change will affect all types of workers with the only exception that the effect of an increase in A_M is immediate both on I_L and I_H.
In this model, an increase in A_H, for example, will increase the wage for high skill workers while it could lower the wage of medium skill workers that are displaced from some of the tasks they previously performed.

Introducing automation in the generalized model presented above requires to provide some rules defining automatable and automated tasks. *Automatable tasks* are all tasks $i \in I \equiv [I', I'']$. Their assumption, based on the findings of Autor et al., 2003, is that this set of automatable tasks must be a subset of the tasks usually carried out by middle skill workers (i.e. $[I', I''] \subset [I_H, I_L]$). However, this assumption can be changed quite easily without affecting the complexity of the model. The set I is exogenously defined as a set of tasks where capital (or machines) has suddenly a comparative advantage over every type of human labour force. For any other task $i \notin I$ capital is considered to be useless in production so that $\alpha_K(i) = 0$.
Automated tasks correspond in this model to automatable task. Due to the fact that technological advancement is exogenous and introduced only in a comparative static analysis.

Following the example presented in the paper, when automation displaces middle skill workers, the model predicts that the ratio w_H/w_M increases, w_M/w_L decreases, while the behavior of w_H/w_L depends on the initial levels of I_L and I_H.

2.1.3 Aghion-Jones-Jones (2017) model

In their working paper, Aghion et al., 2017 present another readaptation of Zeira, 1998 model. They model is a task-based take to standard neoclassical models. Here the authors use a constant elasticity of substitution

(CES) production function

$$Y_t = A_t \left(\int_0^1 x_{i,t}^{\rho} \, di \right)^{\frac{1}{\rho}}$$

with $\rho < 0$. Note that this production function is quite similar to the one used in Zeira, 1998 with the only exceptions of the CES assumption and the fact that the total factor productivity is no longer a constant. Now $A_t = A_0 e^{gt}$ so that technology grows at a rate g.

The authors do not provide a specific definition of *automatable tasks*.This implies that also in this model there is no clear distinction between automatable and automated tasks.
On the other hand, *automated tasks* are defined as an exogenous fraction β_t of tasks. The model is formally simplified with respect to Zeira, 1998 by the assumption that, once a task has been automated, it is produced exclusively using capital and, viceversa, an automated task is produced only with labour. This assumption allows to rewrite the aggregate production function as follows.

$$Y_t = A_t \left[\beta_t^{\frac{1-\rho}{\rho}} K_t + (1 - \beta_t)^{\frac{1-\rho}{\rho}} L_t \right]$$

In this model, automation (intended as a permanent increase in the schedule $\{\beta_t\}$) has drastically different effects on the labour share based on the assumption made on the elasticity of task substitution and on ρ.
When tasks are complements ($\rho < 0$) automation of one task has simultaneously a labour-augmenting and a capital-depleting effects. When tasks are, instead, substitutes ($\rho > 0$) automation is labour-depleting and capital-augmenting. This is in line with the logic presented in Autor, 2015.

2.1.4 Acemoglu-Restrepo (2018) model

Acemoglu and Restrepo, 2018 present a highly generalized task-based model of growth. Their model can be seen as a generalization of the Acemoglu-Author (2011) model.

The production function of the unique good of the economy in continuous time is

$$Y_t = \tilde{B} \left(\int_{N_t - 1}^{N_t} y_t(i)^{\frac{\sigma-1}{\sigma}} \, di \right)^{\frac{\sigma}{\sigma-1}}$$

The task specific production function is

$$y_t(i) = \begin{cases} \bar{B}(\zeta) \left[\eta^{\frac{1}{\zeta}} q(i)^{\frac{\zeta-1}{\zeta}} + (1-\eta)^{\frac{1}{\zeta}} \left(k(i) + \gamma(i) l(i) \right)^{\frac{\zeta-1}{\zeta}} \right]^{\frac{\zeta}{\zeta-1}}, & i \leq I_t \\ \bar{B}(\zeta) \left[\eta^{\frac{1}{\zeta}} q(i)^{\frac{\zeta-1}{\zeta}} + (1-\eta)^{\frac{1}{\zeta}} \left(\gamma(i) l(i) \right)^{\frac{\zeta-1}{\zeta}} \right]^{\frac{\zeta}{\zeta-1}}, & i > I_t \end{cases}$$

In this formalization, $q(i)$ represents an intermediate good used for production. This term can be interpreted as the technology needed for task production. The marginal cost for $q(i)$ is $\psi > 0$.

Assumption(AR2018) 1 *$\gamma(i)$ is increasing in i.*

Assumption(AR2018) 1 is the same as stating that tasks are orderer from the most labour-intensive to the least. We could think at this as having manual and standardized tasks at the bottom of the distribution and cognitive and empathic tasks are at the top.
Another way to look at the task index i is task automation complexity. Tasks close to $N_t - 1$ are easy for machines while the closer we get to N_t, the more complex tasks become to automate. This is why standardization of a task lowers its index in the task distribution.

This production function introduces two interesting aspects of the model. Firstly, the total factor productivity $\tilde{B}$ is assumed to be constant over time. This implies that technological innovation (and thus economic growth) originates exclusively from automation. Secondly, the set of tasks, while retaining its unit size, can now be changed through the parameter N_t. An increase in N_t corresponds to the introduction of new tasks in the economy. Due to assumption 1, labour initially has a strong comparative advantage over capital in the production of every new task introduced. The introduction of new tasks, however, also changes the lower bound of the task distribution. This can be interpreted as new tasks making some of the existing tasks obsolete or as further standardizing existing tasks, thus lowering their index i associating them to a lower $\gamma(i)$. This interpretation is allowed by the fact that the task space is continuous.

There is a representative household whose goal is to maximize a constant relative risk aversion lifetime utility function with rate of time preference $\rho > 0$ and intertemporal elasticity of substitution $\frac{1}{\theta}$:

$$U(C_t, L_t) = \int_0^\infty e^{-t\rho} u(C_t, L_t)\, dt = \int_0^\infty e^{-t\rho} \frac{\left(Ce^{-\nu(L_t)}\right)^{1-\theta} - 1}{1-\theta}\, dt.$$

where $\nu(L_t)$ measures the disutility from working L_t units of time. The authors assume $\nu(L_t)$ to be continuously differentiable, increasing, and convex and to satisfy $\nu''(L_t) + (\theta - 1) \times (\nu'(L_t))^2 / \theta$ such that $u(C_t, L_t)$ is concave.

In this economy there is a clear distinction between automatable and automated tasks.
In the version of the model with exogenous technologies, *automatable tasks* are clearly defined by the exogenously given thresholds $\{I_t\}_{t=0}^\infty$. Every task $i \in [N_t - 1, I_t]$ can be produced either using labour or capital meaning that it is potentially *automatable*. On the other hand, tasks $i \in (I_t, N_t]$ can be produced only by labour. Keep in mind that in this version of the model, also the schedule for $\{N_t\}_{t=0}^\infty$ is exogenously given.
Automated tasks are instead defined by market prices. There will be an endogenous threshold

$$\tilde{I}_t : \quad \frac{W}{\gamma\left(\tilde{I}_t\right)} = R.$$

This implies that for every task $i \in [N_t - 1, \tilde{I}_t]$ capital is the least expensive production factor and these tasks are *automated*. Instead, for tasks $i \in (\tilde{I}_t, N_t]$ labour has a comparative advantage over capital in task production. The main assumption of the authors is that capital stock K is capped at a level $\bar{K}$ such that $\frac{W_t}{\gamma(N_t)} \leq R_t \,\forall t$.

In equilibrium, there will be a threshold $I_t^* = \min\left\{I_t, \tilde{I}_t\right\}$ such that every task $i \in [N_t - 1, I_t^*]$ is produced with capital. We say that the economy is either *technologically constrained* for $I_t < \tilde{I}_t$, or *resource constrained* for $I_t > \tilde{I}_t$.

To endogenize technological growth, instead, the authors introduce a mass S of scientists that choose either to work on automation $S_I(t)$ or on new task creation $S_N(t)$. Scientist have different productivity levels when working on automation κ_I or on new task creation κ_I. This makes it possible to model the change in *automatable tasks* as

$$\dot{I}_t = \kappa_I S_I(t)$$

and the change in the task space as

$$\dot{N}_t = \kappa_N S_N(t).$$

Automated tasks are always defined by market prices $\{R_t, W_t\}$ and labour productivity $\gamma(i)$.
The authors define the function $n_t = N_t - I_t$ as a proxy for the technological level of the economy. An increase in n_t is caused by new task creation while a decrease by the introduction of new automation technologies.

To encourage the investment in scientific research, the authors introduce a perpetual licence system for intermediates $q(i)$. The technology monopolist holds the intellectual property to $q(i)$ and can produce it at the marginal cost $\mu\psi$ with $\mu \in (0, 1)$. An intermediate can be copied by a fringe of competitors at the marginal cost ψ.
The introduction of a new technology that replaces the intermediate $q(i)$ represents an infringement of the property rights of the monopolist. To introduce its new technology, thus, the innovator must make a take-it-or-leave-it (TOLO) offer to the technology monopolist holding the property rights. This TOLO offer will be such that it repays all the present discounted value of future profits of the monopolist using the inferior capital-intensive (in the case of new tasks introduction) or labour-intensive (in the case of task automation) alternative.

Scientist j will incur in a cost $\chi_I^j Y_t$ when working in task automation and $\chi_N^j Y_t$ when working in new task creation. This implies that, provided wages are $W_I^S(t)$ for automation scientists and $W_N^S(t)$ for new task scientists, scientist j will choose to work in automation only if the benefits are greater than the costs:

$$j \text{ works in automation} \iff \frac{W_I^S(t) - W_N^S(t)}{Y(t)} < \chi_I^j - \chi_N^j.$$

The assumption is that $\chi_I^j - \chi_N^j$ is such that all scientists work at all times.

This model results having different balanced growth paths (BGPs) based on the model parameters.

BGP1) Full automation: in this BGP $n_t = 0$ so that there is full automation of task production.

BGP2) Unique interior BGP: in the case of $n_t^* = n_t = n$ (the economy is technologically constraint), there is a BGP if κ_I/κ_N is greater than an arbitrary level $\bar{\kappa} > 0$. This BGP is unique and globally stable if $\theta = 0$ and unique in the neighbourhood of the BGP and asymptotically stable if $\theta > 0$. In this BGP, the introduction of new tasks allows the economy to maintain a level of employment even with constant automation of tasks.

BGP3) Multiple BGPs: in the case of $n_t^* \neq n_t$ (the economy is resource constraint), there exist multiple locally stable BGPs with κ_I/κ_N in between two arbitrary values $\bar{\kappa} > 0$ and $\underline{\kappa} > 0$ with $\underline{\kappa} < \bar{\kappa}$.

BGP4) No automation: with $n_t^* = 1$ only labour is used in task production with $\kappa_I/\kappa_N < \underline{\kappa}$.

2.1.5 Acemoglu-Restrepo (2024) model

Acemoglu and Restrepo, 2024 present a single-sector, static economic model that can be seen as a special case of the model presented in Acemoglu and Restrepo, 2022. The model is then expanded to a multi-sector version but, for clarity of exposition, this section will focus only on the single-sector framework. The main contribution of this working paper is to provide introduce rent dissipation in the model. When automation targets high-rent tasks, wage losses from automation are amplified by the rent dissipation effect. To align the exposition of this model to the other, I will change the notation of the original model referring to tasks with i instead of x and changing the constant elasticity of task substitution from λ to σ.

The authors consider a mass M of tasks $i \in \mathcal{T}$ that are all complementary on with the others. The production function for the unique good of the economy is the constant elasticity of substitution aggregate of single task productions.

$$y = \left[\frac{1}{M}\int_{\mathcal{T}} (M\, y_x)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}$$

Note that if we set $\mathcal{T} \equiv [N-1, N]$ we have a production function that resembles the one in Acemoglu-Restrepo (2018) model.

A discrete number of worker types indexed by $g \in \mathbb{G} = \{1, 2, \ldots, G\}$ competes with capital in task production. Workers of the same type are endowed with the same productivity on every task $\psi_g(i)$.

$$y_i = \psi_k(i)k(i) + \sum_{g=1}^{G} \psi_g(i) l_g(i)$$

where task specific machines $k_(i)$ are treated as intermediate goods produced at a cost $1/q(i)$. For every *automatable task*, the authors impose $q(i) > 0$. When a task is non-automatable $q(i) = 0$.

The model, then, splits the wage paid by a firm producing task i with labour of a worker of type g in two parts: the actual wage w_g and an exogenous wedge $\mu_g(i)$.

$$w_g(i) = \mu_g(i) w_g$$

Observe that here the (exogenous) wedge is the only component of the wage that is task dependent. Moreover, the assumption taken is that wedges are $\mu_g(i) \geq 1$ so that they either represent a premium that firms have to pay to workers of group g to produce task i, or leaves the group wage unaltered for that task. Additionally, the authors use the normalization of imposing the existence of a positive mass of tasks for which $\mu_g(i) = 1$. This implies that wedges for the same group g in other tasks can be considered in relative terms to those tasks where $\mu_g(i) = 1$.

However, firms will hire workers only at a wage that is no more than the value of the marginal product of labour so that in equilibrium we expect $w_g(i) = \mathrm{VMPL}$. Workers will be assigned randomly to tasks and real group wages w_g will adapt to achieve full employment.

To ensure the existence and uniqueness of an equilibrium, the authors also assume that:

- for each task there is at least one worker's group that can produce it;
- for each group of workers there is at least one task that can be produced only by that group (not even capital can be used);
- comparative advantage is strict;
- the integrals $\int_{\{i:\psi_g(i)>0\}} \psi_g(i)^{\sigma-1}\mu_g(i)^{-\sigma}\, di$ and $\int_{\{i:\psi_g(i)>0\}} \psi_g(i)^{\sigma-1}\mu_g(i)^{1-\sigma}\, di$ are bounded for every $g \in \mathbb{G}$.

Under these conditions, the equilibrium exists and is unique. Such equilibrium is inefficient as it features too little employment in high-rent tasks and inefficient automation of tasks. The latter inefficiency is, perhaps, the most interesting one to analyse.
Considering μ_g as the average wedge component for group g, tasks for which

$$\frac{w_g \cdot \mu_g}{\psi_g(i)} < \frac{1}{\psi_g(i) \cdot q(i)} < \frac{w_g \cdot \mu_g(i)}{\psi_g(i)}$$

are automated but inefficiently. Reallocation some workers to these tasks would yield an increase in output of $p(i)\psi_g(i) - w_g\mu_g$ which is always positive under the previously stated condition.

2.1.6 Final thoughts on the current literature trends

In general, the history of economics of automation can be summarized as a rollercoaster of more pessimistic and more optimistic views alternating one with the other. Back in 1930 John Maynard Keynes recognized the importance of trying to detach from the background noise generated by the contemporaneity of technological development in his influential essay *Economic Possibilities for Our Grandchildren* (Keynes, 1930):

*"We are suffering just now from a bad attack of economic pessimism[...]
I believe that this is a wildly mistaken interpretation
of what is happening to us"*.

Reading back at his words poses one question: is it possible that we are currently suffering from the same *bad attack of economic pessimism*?

It is with this question in mind that I would like to conclude this section with some personal thoughts on the current state of the economics of automation. While scientific and engineering papers are more optimistic towards the effects of AI automation, it seems like in economic studies there exists a similar trend of technological pessimism to the one described by Keynes back in 1930. Current economic theories suggest either that we are going to witness "the end of work"[1] or that "we won't be missed"[2]. Moreover, current economic research on automation is highly focused on labour market and distributional effect while only few sources study the effects of AI in education and trade (Lu and Zhou, 2021). Consider, however, the effects that the introduction of an education system where technological advancement affects the learning process of workers could have on the economics of automation. Current studies highlight how AI could have two contrasting effects on knowledge. On one hand, AI has the potential to revolutionize education with personalized learning programs and the opportunity to truly democratize knowledge with real-time translations and simplified access. On the other hand, some scholars are concerned with a potential "*dumbification*" effect of AI. Numerous studies have shown AI-empowered students have lower incentives to study and report worst results in tests. Similar concerns have been raised in the last decades about other technologies disrupting the way human live their lives like the internet and social media. However, one question remains on this regard: is one effect prevailing on the other? It is currently hard to answer this question through empirical studies due to the recency of the subject. For this reason, in section 2.2 I propose a more theoretical approach to the problem looking at the market effects when the prevailing effect is labour augmenting for typically human tasks.

[1] Rifkin and Heilbroner, 1995.
[2] Restrepo, 2025.

2.2 Automation effects on productivity

In this section we present a modification of the most generalized task-based growth model listed literature excursus made in section 2.1. Starting from the Section 2.1.4, we introduce the ability for automation to alter workers' productivity levels.
The intuition of a potential relationship between automation and workers' productivity derives from empirical observations. Autor, 2015 argues that the effects of automation in the labour markets are linked to the relationship between tasks.
These empirical observations suggest that automation could have two contrasting effects on labour productivity.

1. Task automation could allow workers to focus more on complementary tasks. This case could be interpreted as a problem of time and resource allocation for the workers where automation allows humans to concentrate their time and knowledge in less tasks. In this case automation has a *booster effect* on labour productivity for tasks that are not automated.

2. The automation of a task, however, could also displace workers to more demanding tasks thus requiring a retraining of the workforce. In a model like the one we are considering, where there is a single mass of workers all with the same characteristics, this change can be interpreted as a sudden decrease in labour productivity requiring a retraining to effectively perform new tasks. This effect can be described as a *displacement effect*. Another way to observe at this effect is to consider that automation makes workers less capable to carry out non-automated jobs as they are less trained in performing lower level foundational tasks.

This extension of the original model could result to be useful to investigate the rationale behind the diversities we observe in cross-country labour productivity levels, also within the OECD members[3].

The complexity of introducing this dynamic to the model is the fact

[3] OECD, 2025.

that the automation level I_t^* is an equilibrium object. We overcome this problem considering that there exists a lagged effect of automation on labour productivity τ so that the new productivity function is

$$\Gamma(i|I_{t-\tau}^*, N_t) = A(i|I_{t-\tau}^*, N_t)\gamma(i).$$

This introduces the peculiarity of our model. The automation technology I_t plays also a (lagged) role of labour augmenting technology (in the case most interesting where $A(i|I_{t-\tau}^*, N_t) > 1$).

We follow the same approach as Acemoglu and Restrepo, 2018 in defining a static version of the model with exogenous technological change and exogenous capital stock to then be expanded in more complex and realistic models.

2.2.1 Static model

In the static version of the model we consider the economy to be endowed with two task thresholds $\{I, I_0^*\}$ and one task upper bound N. The *technology threshold* I and *task upper bound* N behave exactly as in the original model. The novelty here is represented by the *equilibrium threshold* I_0^* defining the equilibrium automation level before at the beginning of the economy. The assumption is that $I \geq I_0^*$.
The aggregate production function for the unique final good Y is

$$Y = \tilde{B}\left(\int_{N-1}^{N} y(i)^{\frac{\sigma-1}{\sigma}}\, di\right)^{\frac{\sigma}{\sigma-1}}$$

and the task-specific production function is[4]

$$y(i) = \begin{cases} k(i) + \Gamma(i|I_0^*, N)l(i), & \text{if } i \in [N-1, I] \\ \Gamma(i|I_0^*, N)l(i), & \text{if } i \in (I, N] \end{cases}$$

where $\Gamma(i|I_0^*, N) = A(i|I_0^*, N)\gamma(i)$ is our modified productivity function composed by the original task-specific productivity function $\gamma(i)$ multiplied by a task-specific productivity booster $A(i|I_0^*, N)$ defined by the two exogenous values I_0^* and N. The introduction of this automation productivity booster factor will drastically alter the effects of the model.

Assumption 1 $\gamma(i)$ *is increasing in* i *and* $A(i|I_0^*, N) > 0$ *is such that* $A(i|I_0^*, N) = 1, \forall i \leq I_0^*$.

Under this assumption, $A(i|I_0^*, N)$ is such that only tasks that are not automated at the beginning of the economy can be boosted by technological automation. Moreover, the assumption still allows the tasks-specific productivity booster function either to have a positive effect on productivity when $A(i|I_0^*, N) \geq 1$ or a negative effect for $A(i|I_0^*, N) \in (0, 1]$. For the rest of this thesis, however, we will follow the idea that labour productivity is positively affected by automation and thus impose that $A(i|I_0^*, N) \geq 1$

Of course, the definition of the automation productivity booster function has non-trivial effects on the equilibrium definition. We make assumptions to the shape of this function both considering empirical observation and with the goal to maintaining the characteristics of the original model. Assuming that the labour-augmenting effect of automation prevails over its displacement effect, the productivity booster function must satisfy the

[4]Imposing Assumption 2 of the original paper $\eta \to 0$ for sake of exposition. Results are not affected if the original task-specific production function is used

$$y(i) = \begin{cases} \bar{B}(\zeta)\left[\eta^{\frac{1}{\zeta}} q(i)^{\frac{\zeta-1}{\zeta}} + (1-\eta)^{\frac{1}{\zeta}} \left(k(i) + \Gamma(i|I_0^*, N)l(i)\right)^{\frac{\zeta-1}{\zeta}}\right]^{\frac{\zeta}{\zeta-1}}, & \text{if } i \in [N-1, I] \\ \bar{B}(\zeta)\left[\eta^{\frac{1}{\zeta}} q(i)^{\frac{\zeta-1}{\zeta}} + (1-\eta)^{\frac{1}{\zeta}} \left(\Gamma(i|I_0^*, N)l(i)\right)^{\frac{\zeta-1}{\zeta}}\right]^{\frac{\zeta}{\zeta-1}}, & \text{if } i \in (I, N] \end{cases}$$

imposing assumption two conditions.

following conditions

$$A(i|I_0^*, N) \geq 1, \quad \text{and} \quad \frac{\partial}{\partial N} A(i|I_0^*, N) < 0.$$

Moreover, we assume that task automation has a stronger booster effect for higher level tasks relatively to the tasks close to I_0^*, thus requiring

$$\frac{\partial}{\partial i} A(i|I_0^*, N) \geq 0, \quad \text{and} \quad \frac{\partial}{\partial I_0^*} A(i|I_0^*, N) > 0 \text{ only for some } i > \bar{I}$$

where the value for $\bar{I}$ will be defined in detail later.
This implies that workers whose tasks have been automated will use more of their time to produce tasks they were already better at, rather than tasks they were not so good at. This assumption will be key to determine the equilibrium results of our model. Under these assumptions, we propose a Gaussian functional form for the automation productivity booster function where the mean is set to correspond with the highest task $\mu_A = N$ and the variance $\sigma_A = \frac{N - I_0^*}{5I_0^* - 5N + 8}$ is such that the boost is close to negligible for tasks near to I_0^*.

$$A(i|I_0^*, N) = \begin{cases} 1, & i \in [N-1, I_0^*] \\ 1 + \frac{1}{\sqrt{2\pi\sigma_A^2}} e^{-\frac{(i-\mu_A)^2}{\sigma_A^2}}, & i \in (I_0^*, N] \end{cases} \tag{2.3}$$

In the special case where the economy starts with full automation so that $I_0^* = I^* = N$, we impose $A(i|I_0^*, N) = 1$ for every task $i \neq N$ and $A(N|I_0^*, N) = \infty$. This reflects the limit behaviour of the productivity booster function presented in Equation 2.3. This behaviour will be central to our analysis of the model.
Figure 2.3 compares our productivity function with the one used in the original model. A first observation might be that labour productivity for task N has drastically increased due to the automation productivity boost effect. On the other hand, tasks that where previously already automated present exactly the same productivity level. This change alters the comparative advantage schedule of labour over capital attributing to labour a relatively stronger comparative advantage for tasks close to N relatively other tasks.

Imposing a task-specific productivity booster function with this formulation allows us to introduce some interesting behaviours in the model. Figure 2.1 depicts the change in productivity due to an increase in the initial level of automation. The Gaussian shape of the productivity booster function allows us to model the effect of automation making human labour less productive in some middle tasks while increasing productivity for higher (more human) tasks. More precisely, the change from I_0^* to a higher initial automation threshold $I_0^{*\prime}$ will cause:

- no labour productivity change for tasks $i \in [N-1, I_0^*]$;
- a decrease in labour productivity for all tasks $i \in (I_0^*, \bar{I})$;
- no productivity change for a threshold task $\bar{I} \in (I_0^{*\prime}, N)$;
- an increase in labour productivity for all tasks $i \in (\bar{I}, N]$.

Note that the threshold task $\bar{I}$, here defined as the only task higher than I_0^* where labour productivity is uneffected by the change in initial automation, is itself a function of the two values I_0^* and $I_0^{*\prime}$ as well as N.
On the other hand, we expect the introduction of new tasks to cause a decrease in the task-specific productivity for all the tasks that were performed by labour at the beginning of the economy. Tasks $i < I_0^*$ are not affected by the change. This is exactly the behaviour that we observe in Figure 2.2.

A representative household wants to maximize its utility

$$u(C, L) = \frac{\left(Ce^{-\nu(L)}\right)^{1-\theta} - 1}{1-\theta} \tag{2.4}$$

where C is consumption, L is labour supply, $\nu(L)$ is a continuously differentiable, increasing, and convex function representing the utility cost of labour supply, and θ will represent the inverse of the intertemporal elasticity of substitution in the dynamic extension of the model. $\nu(L)$ must satisfy $\nu''(L) + (\theta - 1) \times (\nu'(L))^2 / \theta > 0$ to ensure that $u(C, L)$ is concave.

Assumption 2 *Capital stock K is capped at a level $\bar{K}$: $R = \frac{W}{\Gamma(N|I_0^*, N)}$.*

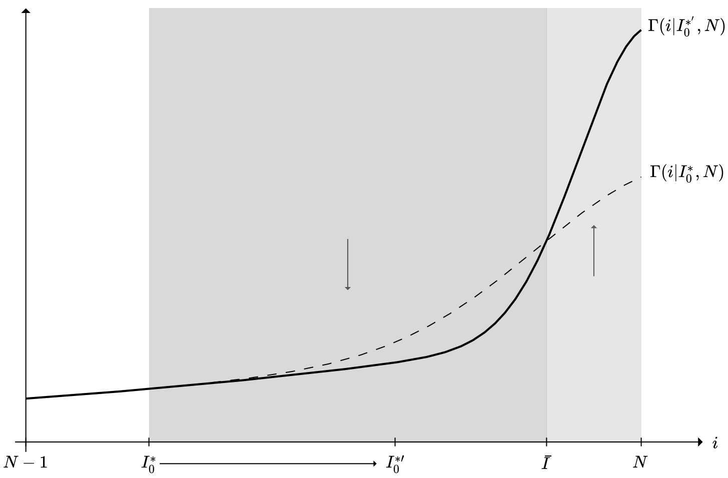

FIGURE 2.1: EFFECT ON TASK-SPECIFIC LABOUR PRODUCTIVITY AFTER AN INCREASE IN THE INITIAL LEVEL OF AUTOMATION

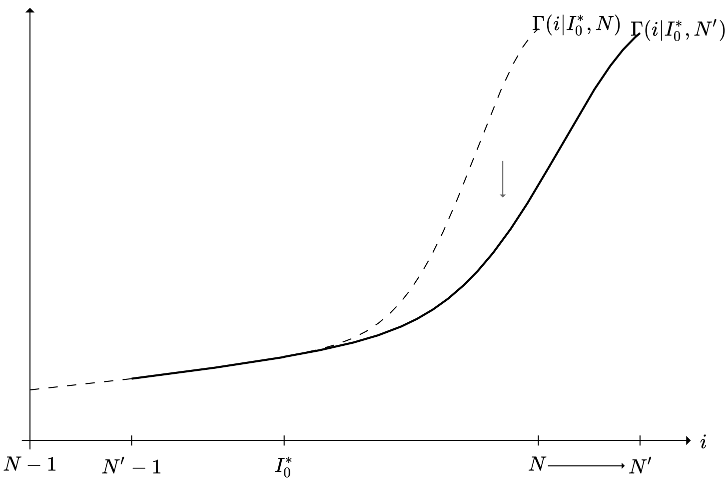

FIGURE 2.2: EFFECT ON TASK-SPECIFIC LABOUR PRODUCTIVITY AFTER THE INTRODUCTION OF NEW TASKS

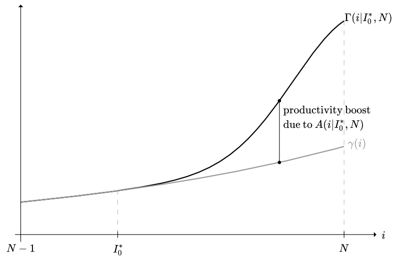

FIGURE 2.3: COMPARING THE PRODUCTIVITY FUNCTION OF THE ORIGINAL MODEL (GREY) AND THE BOOSTED PRODUCTIVITY FUNCTION (BLACK).

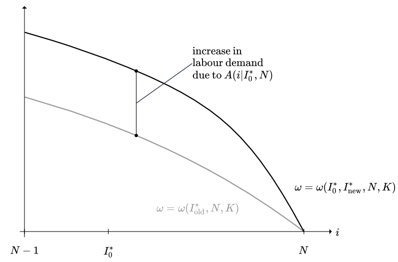

FIGURE 2.4: COMPARING THE RELATIVE LABOUR DEMAND OF THE ORIGINAL MODEL (GREY) AND THE MODIFIED VERSION OF THE MODEL (BLACK).

This corresponds to assumption 3 of the original paper adapted using the new labour productivity function. It ensures that producing new tasks with capital is always inefficient ensuring that new tasks will be always immediately adopted. This assumption will hold only for the static version of the model in which capital stock is exogenous. In Section 2.2.2 the dynamic model will endogenized capital accumulation through household saving decisions.

STATIC EQUILIBRIUM

Given the initial automation level I_0^*, the current technological levels I and N, and the capital stock K, a *static equilibrium* is an allocation $\{W, R, \tilde{I}, I^*, L\}, Y$ of factor prices $\{W, R\}$, threshold tasks $\tilde{I}$ and I^*, employment level L, and aggregate output Y such that all the following conditions hold.

Eq. 1) Task thresholds $\tilde{I}$ and I^* are optimal.

$$\tilde{I}: \; R = \frac{W}{\Gamma(\tilde{I}|I_0^*, N)} \quad \text{and} \quad I^* = \min\{I, \tilde{I}\} \tag{2.5}$$

If we assign to the labour productivity function an exponential form such that $\gamma(i) = e^{\alpha i}$ with $\alpha > 0$, then the value of $\tilde{I}$ can be computed.

$$\Gamma(\tilde{I}|I_0^*, N) = \frac{W}{R}$$

This equation solves for two values of $\tilde{I}$ but only one of them is in the task space $\tilde{I} \in [N-1, N]$ and, thus, is an acceptable solution. This proves existence and uniqueness fo both equilibrium task thresholds.

Eq. 2) The capital and labour markets clear.

$$\begin{aligned} &B^{\sigma-1}(1-\eta)YR^{-\sigma}(I^*-N+1) = K \\ &B^{\sigma-1}(1-\eta)Y\int_{I^*}^{N}\frac{1}{\Gamma(i|I_0^*,N)}\left(\frac{W}{\Gamma(i|I_0^*,N)}\right)^{-\sigma} di = L \end{aligned} \tag{2.6}$$

Eq. 3) The ideal price index condition holds.

$$(I^*-N+1)R^{1-\sigma} + \int_{I^*}^{N}\left(\frac{W}{\Gamma(i|I_0^*,N)}\right)^{1-\sigma} = B^{1-\sigma} \tag{2.7}$$

Eq. 4) The labour supply function can be expressed as an increasing function

$$L = L^S\left(\frac{W}{RK}\right). \tag{2.8}$$

From the representative household utility maximization we get $\nu'(L) = \frac{W}{C}$. Given the static environment and market clearing assumptions, consumption must satisfy $C = RK + WL$.

If we suppose Assumptions 1 and 2 hold, then a static equilibrium exists and is unique. In this equilibrium, the total output is

$$Y = \frac{B}{1-\eta}\left[(I^*-N+1)^{\frac{1}{\sigma}}K^{\frac{\sigma-1}{\sigma}} + \left(\int_{I^*}^{N}\Gamma(i|I_0^*,N)^{\sigma-1}\,di\right)^{\frac{1}{\sigma}}L^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}} \tag{2.9}$$

The equilibrium can be graphically represented by the intersection of two curves in the (i,ω) space where ω is the wage level normalized by capital income

$$\omega = \frac{W}{RK}$$

In this space there exists an increasing curve representing the optimal allocation of production factors across tasks (the *resource curve*) and a decreasing curve plotting the labour demand for every possible level of equilibrium automation (the *demand curve*).

The resource curve can be obtained substituting ω in Equation 2.5.

$$\Gamma(\tilde{I}|I_0^*, N) = \omega K$$

The shape of this curve depends on the labour productivity function Γ and on the (exogenous) capital stock level K. Differently to the original paper, however, this equation now accounts also for the automation productivity boost. This implies that the curve will be higher for tasks not automated at the beginning of the economy.
The demand curve, instead, can be obtained from the market clearing conditions Equation 2.6 and the ideal price index condition Equation 2.7. Rearranging these equilibrium conditions implicitly defines the demand curve with the following equation.

$$\ln(\omega) + \frac{1}{\sigma}\ln(L^S(\omega)) = \left(\frac{1}{\sigma} - 1\right)\ln(K) + \frac{1}{\sigma}\ln\left(\frac{\int_{I^*}^{N} \Gamma(i|I_0^*, N)^{1-\sigma}\, di}{I^* - N + 1}\right)$$

This curve is affected by changes in the initial level of automation I_0^*, the equilibrium automation threshold I^*, the last task N, and the capital stock K. For this reason we will refer to it with the notation $\omega = \omega(I_0^*, I^*, N, K)$.

Figure 2.5 allows us to observe the behaviour of the demand curve when an exogenous increase in I_0^* occurs. When the initial level of automation, the increase in labour productivity causes the demand for labour in the few human tasks left to increase. However, it must be noted that this increase in not homogeneous across tasks For what regards changes in the other variables, the demand curve responds exactly in the same way as in the original model. The introduction of new tasks homogeneously increases the demand curve in all tasks (effectively shifting it to the right). An increase in capital stock intuitively makes capital less expensive and lowers the labour demand curve. The behaviour of the resource curve, instead, is entirely determined by changes in the labour productivity function that we previously discussed while introducing the task-specific productivity booster function.

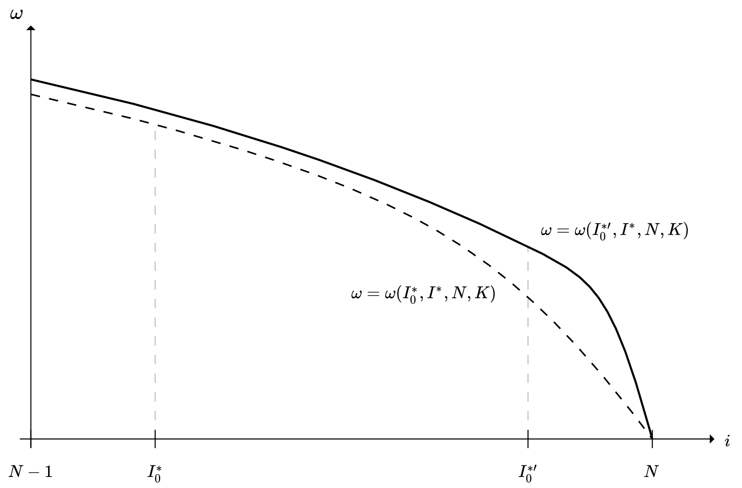

FIGURE 2.5: CHANGE IN LABOUR DEMAND AFTER AN INCREASE IN THE INITIAL LEVEL OF AUTOMATION

STATIC EQUILIBRIA COMPARED TO THE ORIGINAL MODEL

One could wonder what is the effect of the task-specific productivity booster function introduced in the model on the definition of a static equilibrium and how these effects are able to better capture the relationship between automation and labour productivity. Note that the introduction of the productivity booster function $A(i|I_0^*, N)$ has contrasting effects on equilibrium conditions.

The first direct effect depicted in Figure 2.3 is the increase in the threshold $\tilde{I}$ defining the cost-effective allocation of tasks to production factors. The boost in productivity makes labour more competitive in unautomated tasks, thus raising the level of employment associated with a particular value of $\tilde{I}$ compared to the original model.

A secondary effect depicted in Figure 2.4, is the shift of the relative labour demand curve upwards. The new relative labour demand curve accounts also for the productivity boost in all the tasks above I_0^*. For this reason, the effect is higher for lower indexed tasks and completely absent for the last task N (full automation) where labour demand remains at zero. To compare equilibrium results of the two models, we introduce the subscript notation old for variables specific to the original model while the levels for the same variables in the modified version of the model are distinguished by the subscript new. Every variable that is presented with none of these subscripts is held constant in the comparison of the two models.

To study how these two effects interact one with the other in equilibrium, we start from the equilibrium factor prices of the original model where

$$W_{\mathrm{old}} = L^{-\frac{1}{\sigma}} B^{\frac{\sigma-1}{\sigma}} Y^{\frac{1}{\sigma}} \left(\int_{I^*_{\mathrm{old}}}^{N} \gamma(i)^{\sigma-1}\, di \right)^{\frac{1}{\sigma}}$$

and

$$R_{\mathrm{old}} = K^{-\frac{1}{\sigma}} B^{\frac{\sigma-1}{\sigma}} Y^{\frac{1}{\sigma}} \left(I^*_{\mathrm{old}} - N + 1 \right)^{\frac{1}{\sigma}}$$

and in our modified model instead equilibrium factor prices are

$$W_{\mathrm{new}} = L^{-\frac{1}{\sigma}} B^{\frac{\sigma-1}{\sigma}} Y^{\frac{1}{\sigma}} \left(\int_{I^*_{\mathrm{new}}}^{N} \Gamma(i|I^*_0, N)^{\sigma-1}\, di \right)^{\frac{1}{\sigma}}$$

and

$$R_{\mathrm{new}} = K^{-\frac{1}{\sigma}} B^{\frac{\sigma-1}{\sigma}} Y^{\frac{1}{\sigma}} \left(I^*_{\mathrm{new}} - N + 1 \right)^{\frac{1}{\sigma}}$$

We compute the relative wage levels normalized by capital income in the two versions of the model as

$$\omega_{\mathrm{old}} = \frac{L^{-\frac{1}{\sigma}}}{K^{\frac{\sigma-1}{\sigma}}} \left(\frac{\int_{I^*_{\mathrm{old}}}^{N} \gamma(i)^{\sigma-1}\, di}{I^*_{\mathrm{old}} - N + 1} \right)^{\frac{1}{\sigma}}$$

and

$$\omega_{\text{new}} = \frac{L^{-\frac{1}{\sigma}}}{K^{\frac{\sigma-1}{\sigma}}} \left(\frac{\int_{I^*_{\text{new}}}^{N} \Gamma(i|I^*_0, N)^{\sigma-1}\, di}{I^*_{\text{new}} - N + 1} \right)^{\frac{1}{\sigma}}$$

and their ratio can be expressed as

$$\frac{\omega_{\text{new}}}{\omega_{\text{old}}} = \left(\frac{I^*_{\text{old}} - N + 1}{I^*_{\text{new}} - N + 1} \cdot \frac{\int_{I^*_{\text{new}}}^{N} \Gamma(i|I^*_0, N)^{\sigma-1}\, di}{\int_{I^*_{\text{old}}}^{N} \gamma(i)^{\sigma-1}\, di} \right)^{\frac{1}{\sigma}} . \tag{2.10}$$

Panel A

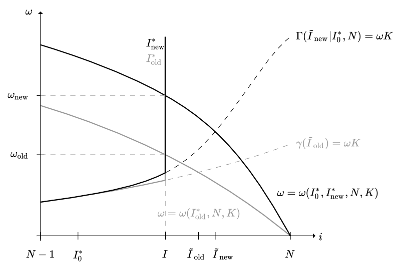

(A) $I < \tilde{I}_{\text{old}} < \tilde{I}_{\text{new}}$

Panel B

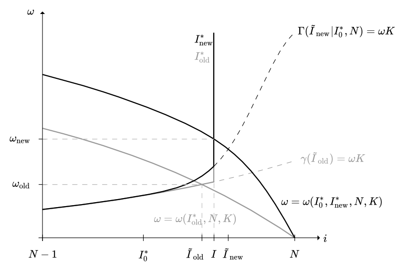

(B) $\tilde{I}_{\text{old}} < I < \tilde{I}_{\text{new}}$

Panel C

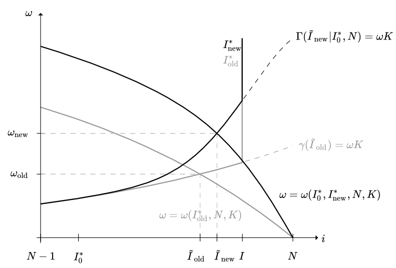

(C) $\tilde{I}_{\text{old}} < \tilde{I}_{\text{new}} < I$

Panel D

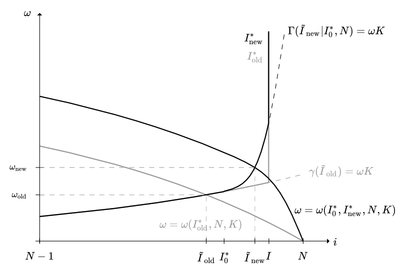

(D) $\tilde{I}_{\text{old}} < I^*_0 < \tilde{I}_{\text{new}} < I$

FIGURE 2.6: DIFFERENCE IN EQUILIBRIUM LEVELS OF LABOUR SHARE BETWEEN THE ORIGINAL MODEL (GRAY) AND OUR MODIFIED VERSION (BLACK).

Figure 2.6 compares the equilibria of the two versions of the model under the same initial conditions. A particular case occurs when the technological threshold I is chosen such that both the original and the modified versions of the model are technologically constraint, meaning that $I^*_{\text{new}} = I^*_{\text{old}} = I$. Under this specific condition, the normalized wage ratio can be rewritten as

$$\frac{\omega_{\text{new}}}{\omega_{\text{old}}} = \left(\frac{\int_I^N \Gamma(i|I_0^*, N)^{\sigma-1}\, di}{\int_I^N \gamma(i)^{\sigma-1}\, di} \right)^{\frac{1}{\sigma}}$$

where the initial assumption $A(i|I_0^*, N) \geq 1$ imposes this ratio to be always greater or equal than one. Panel 2.6A graphically represents this special case. From the plot it is easier to observe that the higher labour supply in the modified model is due exclusively to the increase in the labour demand curve. Being automation fixed in both economies, the only effect that can alter the equilibrium is demand for labour and, intuitively, a boost in labour productivity generate a higher demand.

However, when the initial assumptions are changed and the technological threshold I is set at a higher value in the set $(\tilde{I}_{\text{old}}, N]$, the productivity boost plays a crucial equilibrium role also in determining the efficient allocation of production factors to tasks. One could argue that the boost in productivity should efficiently allocate more tasks to labour, thus lowering the cost efficiency threshold leading to $\tilde{I}_{\text{old}} > \tilde{I}_{\text{new}}$. However, this efficiency effect is contrasted by the labour demand effect in equilibrium. Observing Figure 2.6 hints that the inverse relationship $\tilde{I}_{\text{old}} < \tilde{I}_{\text{new}}$ should always hold. This result could seam counterintuitive at first, but it is a natural consequence of the assumptions taken to modify the model. Consider that the equilibrium now is altered also by the thresholds I_0^*. Recall Figure 2.1 and Figure 2.5 showed the behaviour of both the equilibrium curves when we impose an increase in the initial level of automation. Whilst the resource curve responds differently for tasks above or below the threshold $\bar{I}$, the demand curve always grows. It is the growth in labour demand that ensures that the intuition derived from panels 2.6B, 2.6C, and 2.6D that it must always hold $\tilde{I}_{\text{old}} < \tilde{I}_{\text{new}}$. With this knowledge, we move to a case-by-case analysis of the cases proposed in the three panels.

Panel 2.6B depicts an example where the technological level I is such that the original model is resource constraint while the modified model is technologically constraint. Imposing $I \in \left(\tilde{I}_{\text{old}}, \tilde{I}_{\text{new}}\right)$ also leads the new equilibrium to exhibit higher wages and labour supply. However,

differently from the previous case, now only the modified model is technologically constraint. The equilibrium of the modified model will, thus, exhibit a higher number of tasks allocated to capital compared to the original equilibrium. A similar effect can be observed also in Panel 2.6C where both the economies are resource constraint.
The last panel 2.6D investigates the particular case in which we impose $I_0^* > \tilde{I}_{\text{old}}$. This ensures that the curves $\gamma(i) = \omega K$ and $\Gamma(i|I_0^*, N) = \omega K$ to be the same at $\tilde{I}_{\text{old}}$. Also in this case, the productivity boost $A(i|I_0^*, N)$ will still affect the equilibrium altering the shape of the curve $\omega = \omega(I_0^*, I^*, N, K)$.

The introduction of a boost in productivity caused by an initial level of automation always leads to a higher wage and higher labour supply. However, counterintuitively, it also determines a higher share of automated tasks compared to the original model. This is a direct effect of the choice of imposing the boost in productivity to be bigger for higher index task and lower for lower index tasks. A different assumption might cause also the inverse effect.

Comparative statics

The comparative statics of this model are already drastically different to the ones of the original paper and hint at some of the effects that our change will have on a long run equilibrium. The introduction of the automation productivity booster function effectively creates another variable of interest for comparative statics as well as altering the effects that a change in N has on the equilibrium.

Figure 2.7 shows that, except for level effects, an increase in I has the same general effects on the labour supply as in the original model. Only when the initial level of I is lower than the equilibrium threshold $\tilde{I}$ (Panels 2.7A and 2.7A), an increase in automation technology lowers the labour supply. When instead the economy is resource constraint (Panel 2.7C), advancements in automation technology have no effect on the equilibrium. However, previously we observed how the introduction of new tasks alters also the task-specific productivity. Figure 2.8 shows that an increase in N will both shift the $\Gamma(\tilde{I}|I_0^*, N) = \omega K$ curve downwards and the $\omega(I_0^*, I^*, N, K)$ curve upwards. The effect is an increase in labour supply in all the possible cases (Panels 2.8A, 2.8B, and 2.8C). This is similar to the original model with the only difference that here the effect is exacerbated by the fact that labour productivity decreases. The modified model maintains the main characteristic property of the original model granting to the two technological thresholds I and N polarized effects on employment. On the other hand, it introduces the ability to alter the equilibrium thought another variable: I_0^*.

Panel A

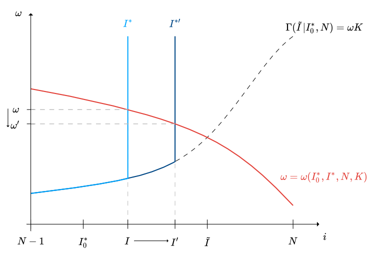

(A) $I < I' < \tilde{I}$

Panel B

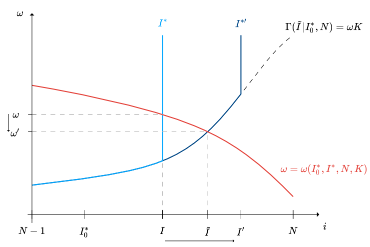

(B) $I < \tilde{I} < I'$

Panel C

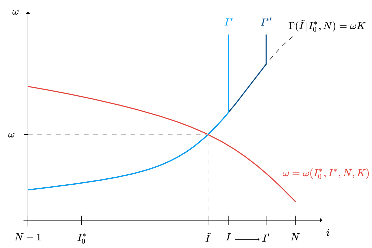

(C) $\tilde{I} < I < I'$

FIGURE 2.7: CHANGES IN NORMALIZED WAGE WITH THE INTRODUCTION OF NEW AUTOMATION TECHNOLOGIES.

In fact, the novelty now is represented by the fact that the equilibrium can be affected also by the level of initial automation I_0^*. This variable allows for two interesting comparative static analysis. As first, Figure 2.9 shows two economies with different initial levels of automation $I_0\prime > I_0^*$ but keeps their current technological levels I and N constant. In general, we observe that the economy with the lower initial automation level I_0^* has also a lower labour share (Panels 2.9A and 2.9B). In the case in which the

Panel A

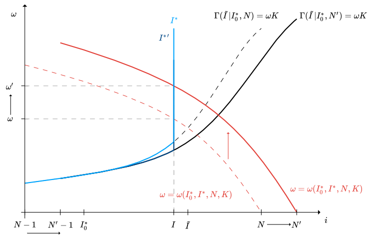

(A) $I < \tilde{I} < \tilde{I}'$

Panel B

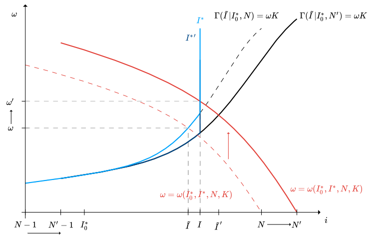

(B) $\tilde{I} < I < \tilde{I}'$

Panel C

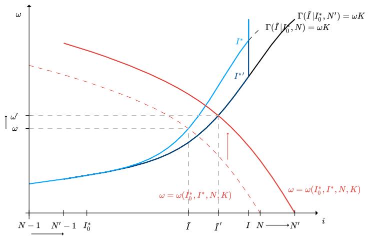

(C) $\tilde{I} < \tilde{I}' < I$

Figure 2.8: Changes in normalized wage with the introduction of new tasks.

increase in I_0^* is such that $\tilde{I} < \tilde{I}' < I$ so that both economies are resource constraint (Panel 2.9C), the relationship between ω' and ω depends on the magnitude of the change in I_0^*. If the change is too small it might occur that ω' is marginally bigger than ω, but when the ΔI_0^* is high enough, the relationship inverts and $\omega' < \omega$. This suggests that there must exist a ΔI_0^* such that $\omega' = \omega$ for every possible value of I_0^*. This analysis, however, assumes the economy with the lower level of initial automation I_0^* will increase its automation technology more than the other economy. In particular, if the difference in the initial level of automation between the two economies is $\Delta I_0^* > 0$, the change in technological automatability for the first economy is

$$I - I_0^* = I - I_0^{*\prime} + \Delta I_0^* > I - I_0^{*\prime}$$

A secondary approach could be to consider also a change in technological automation I such that $I' = I + \Delta I_0^*$. The upper bound of the task distribution N is kept constant. This is the analysis shown in Figure 2.10. When both economies are technologically constraint (Panel 2.10A) the downward shape of the $\omega(I_0^*, I, N, K)$ curve causes the economy with a higher technological automation to have a lower labour share. When both economies are resource constraint, instead, the change in labour share due to an a higher technological level is ambiguous.

The introduction of the productivity booster function $A(i|I_0^*, N)$ allows us to model the observation that workers naturally adapt to automation specializing in more *human tasks*. This allows workers to maintain some occupation also with near to full automation. The I_0^* term is now exogenously give but, provided that the results of the model heavily rely on its value, it will be endogenized in **??**.

2.2.2 Dynamic model

This section extends the previous static model to a dynamic environment in which the evolution of the capital stock is determined by the rational

Panel A

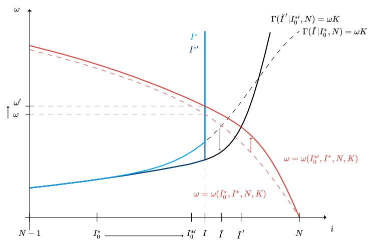

(A) $I < \tilde{I} < \tilde{I}'$

Panel B

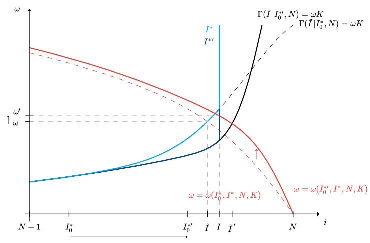

(B) $\tilde{I} < I < \tilde{I}'$

Panel C

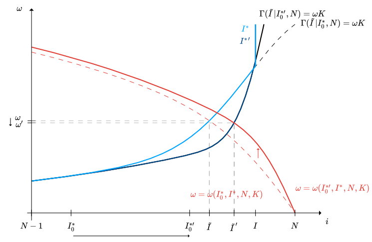

(C) $\tilde{I} < \tilde{I}' < I$

FIGURE 2.9: CHANGES IN NORMALIZED WAGE IN TWO ECONOMIES WITH DIFFERENT INITIAL LEVELS OF AUTOMATION.

Panel A

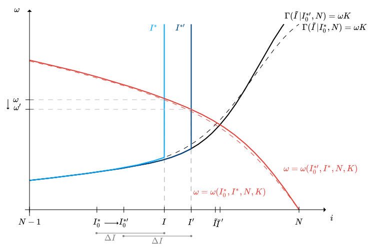

(A) $I < \tilde{I}, \quad I' < \tilde{I}'$

Panel B

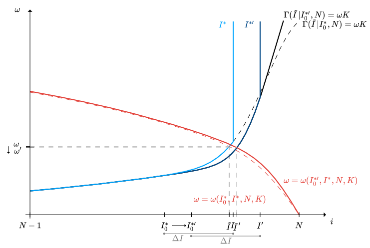

(B) $I > \tilde{I}, \quad I' > \tilde{I}'$

FIGURE 2.10: CHANGES IN NORMALIZED WAGE IN TWO ECONOMIES WITH DIFFERENT INITIAL LEVELS OF AUTOMATION CONTROLLING FOR AUTOMATION GROWTH.

choices of a representative household maximizing their lifetime utility

$$\int_0^\infty e^{-\rho t} u(C_t, L_t)\, dt$$

where $u(C_t, L_t)$ is as defined in Equation 2.4 and $\rho > 0$ is the rate of time preference. Capital depreciates at a rate δ.

The production function of the unique final good is

$$Y_t = \tilde{B} \left(\int_{N_t - 1}^{N_t} y_t(i)^{\frac{\sigma-1}{\sigma}}\, di \right)^{\frac{\sigma}{\sigma-1}}$$

and the task-specific production function is[5]

$$y_t(i) = \begin{cases} k_t(i) + \Gamma(i|I^*_{t-\tau}, N_t) l_t(i), & \text{if } i \in [N_t - 1, I_t] \\ \Gamma(i|I^*_{t-\tau}, N_t) l_t(i), & \text{if } i \in (I_t, N_t] \end{cases}$$

where we change the labour productivity booster function from the static model to

$$A(i|I^*_{t-\tau}, N_t) = \begin{cases} 1, & i \in [N_t - 1, I^*_{t-\tau}] \\ 1 + \frac{1}{\sqrt{2\pi\sigma_A^2}} e^{-\frac{(i-\mu_A)^2}{\sigma_A^2}}, & i \in (I^*_{t-\tau}, N] \end{cases}$$

with $\mu_A = N_t$ and $\sigma_A = \frac{N_t - I^*_{t-\tau}}{5I^*_{t-\tau} - 5N_t + 8}$ where $\tau \geq 0$. With this change, $A(i|I^*_{t-\tau}, N)$ now represents a lagged effect of technological automation on the task-specific labour productivity. The effect of new task creation is instead considered to be instantaneous. Following the same assumptions of the original model, new tasks are immediately produced the same period as they are invented. Note that the productivity booster function cannot be defined for $I^*_{t-\tau} = N_t$. We assume that in such a case $A(i|I^*_{t-\tau}, N_t) \to \infty$.

We follow the original paper in summarizing the technological level of the economy as

$$n_t = N_t - I_t$$

[5] Imposing Assumption 2 of the original paper $\eta \to 0$ for sake of exposition. Results are not affected if the original task-specific production function is imposed.

and the state of technology used in equilibrium as

$$n_t^* = N_t - I_t^*$$

However, with our change in the comparative advantage schedule, the technological level of the economy is now affected also by the level of previous automation $I_{t-\tau}^*$. The assumption is that technological innovation can only grow over time so that it must always hold that $I_t \geq I_{t-\tau}^*$. n_t and n_t^* can still be considered a proxy of the technological level of the economy at time t.

Equation 2.9 from the static version of the model can be rewritten here like

$$Y_t = B\left[(I_t^* - N_t + 1)^{\frac{1}{\sigma}} K_t^{\frac{\sigma-1}{\sigma}} + \left(\int_{I_t^*}^{N_t} \Gamma(i|I_{t-\tau}^*, N_t)^{\sigma-1}\, di\right)^{\frac{1}{\sigma}} L_t^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}$$

which can be rewritten as a function of the technological level n_t^*, the equilibrium productivity $\Gamma(I_t^*|I_{t-\tau}^*, N_t)$, the capital stock K_t, the employment level L_t as well as the technological thresholds I_t^*, $I_{t-\tau}^*$, and N_t. To simplify the notation we define two functions $a_K = (1 - n_t^*)^{\frac{1}{\sigma}}$ and $a_L = \left[\int_0^{n_t^*} \left(\gamma(i) e^{-\frac{i^2 - 2(N_t + I_t^*)i}{\sigma_A^2}}\right)^{\sigma-1} di\right]^{\frac{1}{\sigma}}$ so that output can be written as

$$Y_t = \left\{a_k K_t^{\frac{\sigma-1}{\sigma}} + a_L \left(\Gamma(I_t^*|I_{t-\tau}^*, N_t) L_t\right)^{\frac{\sigma-1}{\sigma}}\right\}^{\frac{\sigma}{\sigma-1}}$$

With this formulation, factor prices can be easily determined in terms of the employment level L_t and the other capital, consumption, and wages normalized by labour's equilibrium comparative advantage

$$k_t = \frac{K_t}{\Gamma(I_t^*|I_{t-\tau}^*, N_t)}, \quad c_t = \frac{C_t e^{\frac{1-\theta}{\theta}\nu(L_t)}}{\Gamma(I_t^*|I_{t-\tau}^*, N_t)}, \quad \text{and} \quad w_t = \frac{W_t}{\Gamma(I_t^*|I_{t-\tau}^*, N_t)}.$$

Factor prices can now be expressed as

$$R_t = B\,a_K \left[a_K + a_L \left(\frac{L_t}{k_t} \right)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{1}{\sigma-1}}$$

and

$$w_t = B\,a_L \left[a_K \left(\frac{k_t}{L_t} \right)^{\frac{\sigma-1}{\sigma}} + a_L \right]^{\frac{1}{\sigma-1}}$$

where $\frac{k_t}{L_t}$ essentially represents capital in effective units of labour.

DYNAMIC EQUILIBRIUM

Let g_t be the growth rate of $\Gamma(I_t^* | I_{t-\tau}^*, N_t)$ at time t. Known the levels of the growth rate g_t, the technological threshold n_t, and the delayed automation factor $I_{t-\tau}^*$, a *dynamic equilibrium* can be defined as a path for the threshold n_t^* as well as normalized consumption, normalized capital, and employment $\{c_t, k_t, L_t\}$ that satisfies:

Eq. 1) $n_t^* \geq n_t$ and, in particular,

$$n_t^* = n_t \Rightarrow w_t > R_t \quad \wedge \quad n_t^* > n_t \Rightarrow w_t = R_t;$$

Eq. 2) the Euler equation

$$\frac{\dot{c}_t}{c_t} = \frac{1}{\theta}(R_t - \delta - \rho) - g_t;$$

Eq. 3) the endogenous labour supply condition stating that labour income must equate the marginal disutility from labour in equilibrium

$$\nu'(L_t) e^{\frac{\theta-1}{\theta}\nu(L_t)} = \frac{W_t}{c_t};$$

Eq. 4) the dynamic resource constraint

$$\dot{k}_t = R_t - c_t e^{-\frac{1-\theta}{\theta}\nu(L_t)} - (\delta - g_t) k_t;$$

Eq. 5) the transversality condition requiring all capital to be consumed by the end of the economy

$$\lim_{t\to\infty} k_t e^{-\int_0^t R_s-\delta-g_s\,ds} = 0$$

With respect to the original model, now the economy no longer admits a full automation equilibrium. Due to the introduction of the productivity booster function, full automation would now require capital to infinitely grow in equilibrium but the transversality condition does not allow such behaviour. This exclusion of a full automation equilibrium is a direct effect of the functional form we attributed to the productivity booster function.

Balanced Growth Paths

In this economy, a *balanced growth path* (BGP) can be defined as a dynamic equilibrium where the economy grows at a constant positive rate g, factor shares are constant, and the rental rate of capital R_t is constant.

Factor shares are constant for a constant value of the technological spread $n_t^* = n^*$.

In equilibrium, the economy will grow at the same rate g_t as the labour productivity at the highest task produced with capital $\Gamma(I_t^*|I_{t-\tau}^*, N_t)$. Adopting the notation of using a dotted variable to indicate its derivative with respect to time, the growth rate of total output can be written $\dot{\Gamma}(I_t^*|I_{t-\tau}^*, N_t)/\Gamma(I_t^*|I_{t-\tau}^*, N_t)$. Given that $\Gamma(i|I_{t-\tau}^*, N_t)$ is defined as the product between two functions, we can write

$$g_t = \frac{\dot{\Gamma}(I_t^*|I_{t-\tau}^*, N_t)}{\Gamma(I_t^*|I_{t-\tau}^*, N_t)} = \frac{\dot{A}(I_t^*|I_{t-\tau}^*, N_t)}{A(I_t^*|I_{t-\tau}^*, N_t)} + \frac{\dot{\gamma}(I_t^*)}{\gamma(I_t^*)}$$

Both growth rates are always non-negative in the task space $[N-1, N]$, independently of the evolution of N_t, and I_t^*.
The only difference with the original model is that, to impose constant positive growth, it is necessary to introduce one extra condition. Necessary and sufficient conditions for a BGP are $n_t^* = n^*$ and $I_t^* - I_{t-\tau}^* = \Delta_\tau$. The latter is an ex-ante condition required to achieve balanced growth in

this economy, while in the original model the same condition is imposed ex-post as a consequence of sustained growth.

$n_t^* = n^*$ and $I_t^* - I_{t-\tau}^* = \Delta_\tau$ are sufficient BGP conditions since they lead the productivity booster function to be constant over time $g_A = 0$ while the growth rate of the comparative advantage schedule to be positive $g_\gamma = \alpha\frac{\Delta_\tau}{\tau} > 0$ and constant. The economy will thus grow at a constant positive rate $g = g_t = g_A + g_\gamma = \alpha\frac{\Delta_\tau}{\tau}$.

Proving that $n_t^* = n^*$ and $I_t^* - I_{t-\tau}^* = \Delta_\tau$ are necessary BGP conditions can, instead, be more complex. From the market clearing conditions we get that

$$R_t^\sigma \frac{K_t}{Y_t} = B^{\sigma-1}(1 - n_t^*)$$

and

$$W_t^\sigma \frac{L_t}{Y_t} = B^{\sigma-1}\Gamma(I_t^*|I_{t-\tau}^*, N_t)^{\sigma-1} a_L^\sigma$$

while the ideal price index condition imposes

$$(1 - n_t^*)R_t^{1-\sigma} + \frac{W_t}{\Gamma(I_t^*|I_{t-\tau}^*, N_t)^{\sigma-1}} a_L^\sigma = B^{1-\sigma}$$

From the first equation, since both R_t and K_t/Y_t must be constant in a BGP, we obtain that also n_t^* must be constant. To satisfy BGP condition, then, the other two equations require $\Gamma(I_t^*|I_{t-\tau}^*, N_t)$ to stay constant over time which happens when the inter-period automation change $I_t^* - I_{t-\tau}^* = \Delta_\tau$ is constant under $n_t^* = n^*$.

Under these two conditions, the economy has similar BGPs as the original model with the only exception being that we have to discard the full-automation BGP given it does not satisfy equilibrium conditions. The economy has three different of BGPs:

BGP 1) Interior BGP with Immediate Automation: in this BGP $I_t^* = I_t \, \forall t$ so that $n_t^* = n_t = n \in (0, 1)$. The second condition imposes $\dot{I}_t^* = \Delta$, and since $n_t^* = n$ this requires also $\dot{N}_t = \dot{I}_t = \dot{I}_t^* = \Delta$. The growth rate of aggregate output is $g = \alpha\Delta$. These conditions identify a unique and globally stable BGP.

BGP 2) Interior BGP with Eventual Automation: this BGP imposes $I_t^* = \tilde{I}_t < I_t \, \forall t$ so that $n_t^* = n^* \in (0, 1)$ and $n_t < n^*$. Again, it must hold that $\dot{I}_t^* = \dot{N}_t = \Delta$ and the growth rate of the economy is give by $g = \alpha\Delta$ but in this case $\dot{I}_t \leq \Delta$. This BGP is unique and globally stable.

BGP 3) No Automation: tasks can be produced only with labour and still lead to balanced growth. In this unique and globally stable BGP $n_t^* = 1$ and $n_t \geq 1$, meaning that the capital share is 0. Sustained growth still requires $\dot{I}_t^* = \dot{N}_t = \Delta$ which is a trivial condition considering that this BGP imposes $I_t^* = N_t - 1$.

Proofs follow an identical approach to the one used in Appendix A of the original paper.

Long-run Comparative Statics

This section studies the effects of an unanticipated and permanent decrease in n_t on the endogenous variables starting from an interior balanced growth path of type 1. Recall that a BGP1 requires the economy to be technologically constraint in equilibrium so that $I_t^* = I_t \, \forall t$. The choice of focusing on a BGP of this kind has been driven by its interesting long-run and short-run effects. Consider that, a decrease in n_t portraits automation running ahead of new task creation. It could be caused both by an sudden increase in automation technology I_T or by an unexpectedly slower growth in new task creation $\dot{N}_T$ at a time T. Differently from the original model, our modification is such that an increase in automation technologies will have slightly different short- and long-term equilibrium effects on endogenous variables.

If the shock in n_t is caused by an unexpected and permanent increase in I_t at a time T, the change will have a lagged effect on productivity also at time $T + \tau$ through the productivity booster function. The lagged effect on productivity, however, will no longer be unanticipated but expected and the economy will smoothly gravitate toward that new equilibrium instead of discontinuously jump from one state to the other. The results of this lagged productivity effect of automation, however, will be different based

on the values of θ and τ. While θ determines how long does it take for households to adapt variables to the new equilibrium, τ instead represents the time period automation needs to affect productivity. A reasonable assumption would be that the effect of automation on productivity occurs before the economy adapts to the new equilibrium conditions. We begin from an analysis of this case to then follow by highlighting the differences arising when the intertemporal elasticity of substitution is sufficiently low or τ is sufficiently high for the economy to reach a new long-run equilibrium before the lagged effect of automation shows its effects.
Figure 2.11 depicts the two different short-term effects that the unanticipated increase in I_t at time T can have on wages. This paragraph focuses on the dashed lines plotting a case where the original value of I_t is too close to $\tilde{I}_t$. The increase in technological automation will move the economy to a resource constraint state. The low productivity gains from the extra automation leads wages to initially decrease in the short run to then smoothly move toward higher wages at time $T + \tau$ due to the change in productivity. It should be considered that, also at time $T + \tau$ two conditions could arise. It could be that the change in labour productivity is such that the economy remains resource constraint (similar to the case presented in Figure 2.9C). In this case (represented by the black dashed line in the plot), we expect a slight decrease in labour supply to cause a small increase in wages. On the other hand, when the change in labour productivity at time $T + \tau$ is such that the economy goes back to be technologically constraint similarly to the case presented in Figure 2.9B, we observe a slight increase in labour supply leading to a small decrease in wages (case represented by the gray dashed line in the plot). These two conditions will present also different long-term effects. While in both cases wages must be higher than the pre-shock wage schedule, when the economy is technologically constraint wages are always higher than in the resource constraint case. The other panels discuss the behaviour of labour supply, capital stock, and rental rate of capital under the same conditions. The labour supply plot in Panel 2.11B depicts the changes argued above. The dashed lines smoothly change to adapt to the expected change at time $T + \tau$. In all three cases, however, labour supply is permanently lower than the pre-shock supply level in the long run. The rental rate of

capital, instead, presents a short-term increase at time T in all three cases followed by different effects at time $T - \tau$. When the economy is resource constraint, we expect the rental rate of capital to be naturally higher as the share of automated tasks is higher. In the long run, capital stock will adjust to bring back the rental rate to the equilibrium level $\delta + \rho + \theta\alpha\Delta$. The different long-run equilibrium levels reflect how much capital must be accumulated to bring back the rental rate to equilibrium.
The assumption that the shock in I_t depicted by the solid line in Figure 2.11 is such that the economy is technologically constraint also after time $T + \tau$ makes it so the solid line does not split into two cases. The only possible equilibrium at time $T + \tau$ is that the economy will still be technologically constraint at the same level I_t. However, at $T + \tau$ there will be an overall increase in labour productivity, causing a slight increase in wages and a decrease in labour supply. The plot purposefully exaggerates the changes in labour supply at time $T + \tau$ for the sake of exposition. However, one should expect this change to be smooth over time with the adaptation period depending on the intertemporal elasticity of substitution. Also the rental rate of capital will be affected by the change in productivity, effectively increasing thanks to the productivity gains. Capital stock will respond to this change effectively lowering capital accumulation to bring back capital prices to the equilibrium level $\delta + \rho + \theta\alpha\Delta$.

When the shock in n_t is instead due to an unanticipated and permanent decrease in N_t, there is no lagged effect to take into consideration after the initial shock at time T. However, the productivity booster function is such that a decrease in N_t increases the productivity of workers in all the tasks not automated at time $T - \tau$ while lowering the labour demand curve as shown in Figure 2.8. This causes a decrease in labour supply and makes capital suddenly more valuable. The short-term effect on wages is similar to the one shown in the case of an unexpected increase in automation with two possible effects depending on the pre-shock level of automation. These effects are the same in both the case that the economy remains technologically constraint or shifts to a resource constraint state after the shock even thought, in the two cases, the shock effects will have different magnitudes as shown in Figure 2.12.

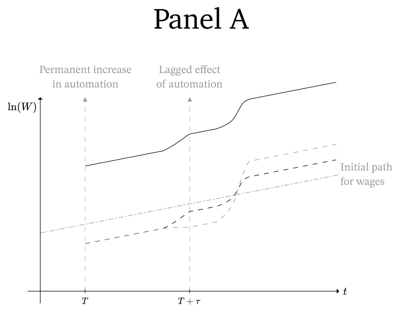

(A) Effects on $\ln(W_t)$

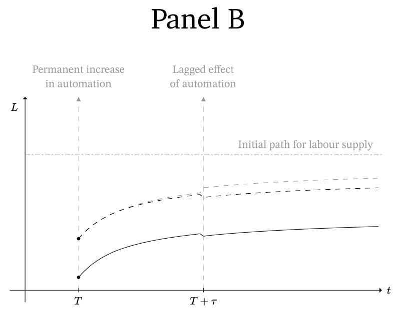

(B) Effects on L_t

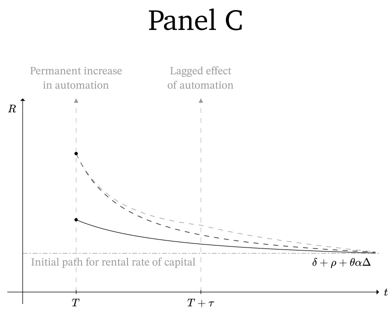

(C) Effects on R_t

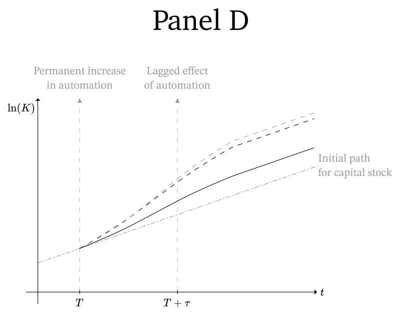

(D) Effects on $\ln(K_t)$

Figure 2.11: Short- and long-run effects of an unanticipated and permanent increase in automation technology I_t starting from a BGP of type 1.

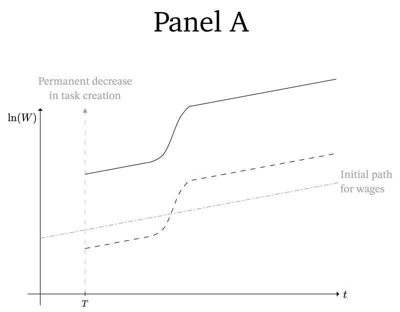

(A) EFFECTS ON $\ln(W_t)$

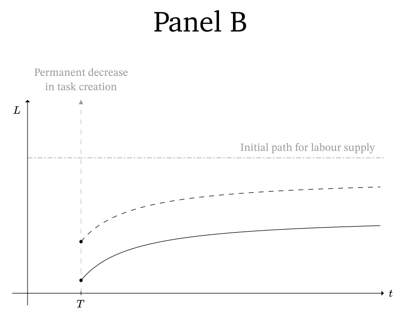

(B) EFFECTS ON L_t

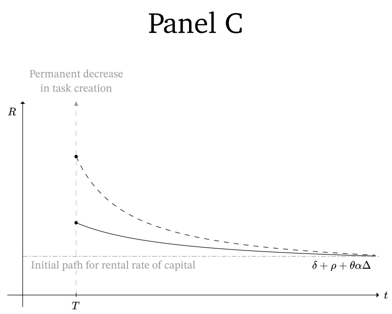

(C) EFFECTS ON R_t

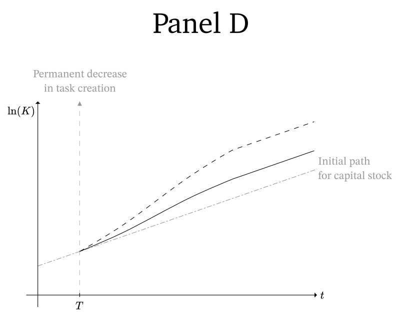

(D) EFFECTS ON $\ln(K_t)$

FIGURE 2.12: SHORT- AND LONG-RUN EFFECTS OF AN UNANTICIPATED AND PERMANENT DECREASE IN TASK CREATION TECHNOLOGY N_t STARTING FROM A BGP OF TYPE 1.

A slightly different case occurs, instead, when θ and τ are such that households converge to the new steady state before the lagged effect of automation on the comparative advantage schedule takes place. The directions of changes will be exactly the same as the ones discussed in the previous case. What changes is the adaptation time of the economy to the new equilibrium. In fact, we expect a first adaptation to a long-run equilibrium determined under the shock conditions to then smoothly gravitate towards a newly changed state at time $T + \tau$ gradually shifting the economy to the final steady state.

LONG-RUN COMPARISON WITH THE ORIGINAL MODEL

Comparing these long-run comparative statics results to the ones reached in the original model is mostly a matter of considering the lagged response of the new model τ units of time after the shock occurred and accounting for the changes in equilibrium factor shares due to the change in the comparative advantage schedule of labour. We expect an increase both the in short-run and in the long-run levels of labour supply and wages under every circumstance going from the original model to the new assumptions. The higher level of equilibrium automation of the modified model also leads to a higher rental rate of capital in case of both the economies being resource constraint after the shock. This requires capital to increase more in the long-run to bring the rental rate at the equilibrium level. These general results should always hold. However, a deeper analysis of the behaviour of the two models would require us to distinguish at least five cases. In the case of an unanticipated and permanent increase in I_t, different behaviours are observed when a combination of the following conditions occurs.

- The pre-shock level of automation technology is
  - sufficiently far away from to the efficiency threshold;
  - too close to the efficiency threshold.
- The lagged effect on productivity is such that
  - both the original model and the new model are technologically

constraint in the long-run;

  - the original model is technologically constraint in the long-run while the modified model occurs to be resource constraint;
  - both the original model and the new model are resource constraint in the long-run.

- The lagged effect on productivity occurs:
  - before the economy reaches the new steady state;
  - after the economy has already reached the new steady state.

Figure 2.13 shows an analysis of the simpler case where both the economies are technologically constrained also after the lagged effect of automation on productivity and the lag parameter τ is small. Moreover, it assumes the pre-shock level of automation technology to be high enough for wages to increase in the short-run as a response to the increase in I_t.

The main results of these comparative statics between the two models resemble the results reached in the comparison of the static versions of the models.

2.2.3 General results and future research

This work began from one empirical observation: human work is always altered by the introduction of new automation technologies. Not only there is an obvious displacement of the workers previously employed in the newly automated tasks, but also there could be a change in the way non-automated tasks are produced by labour. Introducing this behaviour in a growth model can help to us to shape how markets behave to new automation technologies. This attempt at introducing a labour productivity booster function able to alter labour productivity at a task level in the quite flexible model of Acemoglu and Restrepo, 2018 represents a first step towards a more general model. The results of the model proposed in this thesis, in fact, are heavily driven by the initial assumptions on the shape of the productivity booster function. However, the model proposed

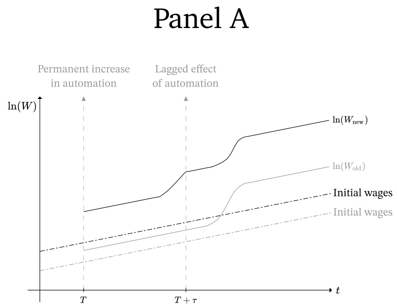

(A) Effects on $\ln(W_t)$

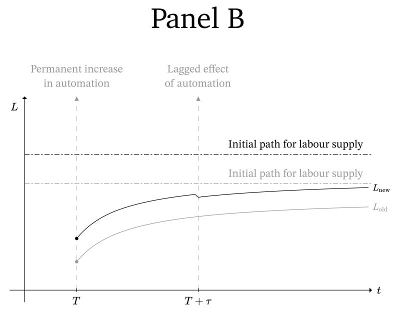

(B) Effects on L_t

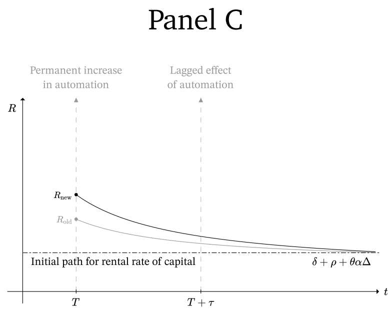

(C) Effects on R_t

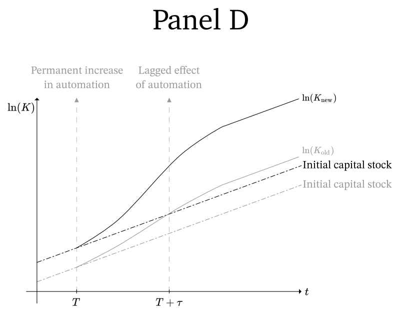

(D) Effects on $\ln(K_t)$

Figure 2.13: Comparison of the modified model long-run response to a shock in n_t compared with the original model

does not lead to trivial results. Once one accounts for positive effects that automation has on human labour (one above all, the ability to use time for more productive, complex, typically human tasks) a full automation equilibrium is immediately excluded. Still, these productivity boosts make also automation more valuable thus raising the efficiency threshold with respect to the non-boosted original model. So, even thought full automation does not represent an equilibrium, there is a tendency to produce more tasks with capital. In addition, this model inherits an important result from Acemoglu and Restrepo, 2018: in equilibrium the growth rates of technological automation and creation of new tasks must not only be constant, but also equal. If the creation of new tasks runs behind technological automation, the economy begins to grow at a pace that is not constant. This condition could seam, at best, far from the current reality where the growth of technological automation appears to be way faster than the introduction of new tasks.

With this, we come to a list of possible extensions of the model presented in this thesis.
At first, there is not enough empirical evidence to prove that automation effect on labour productivity is positive. A required generalization is to allow the productivity booster function to be also lower than 1. However, to maintain Assumption 1 valid it must be that the function is either a productivity booster (greater or equal than one) or a productivity detractor (lower or equal than one) for all tasks. Otherwise, it should be altered the way tasks are indexed to allow both effects to kick in.
Another hypothesis is to tackle the problem from a different perspective. Instead of making labour productivity directly a function of the automation level of the economy, one could assume that a specific automation technology can be used also by households to augment their utility from leisure. This would create a secondary market for the automation technology and introduces a new decision variable for the households' utility maximization problem while allowing workers to allocate more of their time to carry out human tasks.

Chapter 3

AI overtake

The pace at which AI and robotics are advancing nowadays is raising many concerns both from the general public and from filed experts on potential mass automation of entire jobs. The existential question of the last few years became: what would happen to humanity once machines are both autonomous and more capable than humans?

This chapter wants to provide the basis to reflect on this (mostly philosophical) interrogative starting from a vertical focus on the Artificial Intelligence technology in section 3.1 to the look at the first attempts to model an AGI economy in **??**. The model previously developed in chapter 2 attempted a positive theoretical approach at tackling this question based also on the more technical knowledge presented in this chapter.

3.1 A vertical focus on Artificial Intelligence

This section is dedicated to an historical overview of Artificial Intelligence (AI) development (Section 3.1.1) and an extensive research on AI current and future capabilities or limitations (Section 3.1.2). Following the intuition that technology shapes automatability, the discussion to be presented in this section lays down key definition and concepts that will be used thought the whole thesis.

3.1.1 The history of AI

As discussed in the previous section, understanding the technology is essential to define its automation effects. This section will focus on the history of AI as well as the definition of the AI term itself.

Artificial Intelligence is a quite broad term introduced in 1956 by the American physicist John McCarthy as "the science and engineering of making intelligent machines". However, an antilitteram definition of AI can be found in Turing, 1950, considered to be one of the fathers of AI. Struggling to find an answer to the question "Can machines think?", Turing designed the *imitation game* also known as Turing test. The game defines a machine as being intelligent if it is able to answer a series of questions in a way that is indistinguishable from how a human would answer them. In its original form, the test requires a machine to understand texts, be able to formulate and answer and write it in natural language. Further on, scholars proposed an extension of the test to include also other sensory inputs like vision and hearing and/or outputs like synthetic voices.

However, the birth of the idea of an intelligent machine is often attributed to the science fiction writer Isaac Asimov and his book *I Robot* published in 1942 currently used as reference in many AI ethics discussions. More practically speaking, the mathematical foundations for AI development can be found in Turing, 1936 later applied by John von Neumann, Norbert Wiener, and other scholars in the late 1930s and early 1940s. Up until that point, however, AI was still only a theoretical abstraction. I was with McCulloch and Pitts, 1943 that the first mathematical model of an artificial neuron - the *McCulloch-Pitts neuron* - was developed. Since then AI development had a tumultuous path with periods of high (*AI summers*) and lows (*AI winters*).

The first AI summer is often set to begin with the birth of the term AI in 1956 and lasted for nearly two decades. In 1957, advancements like Newell and Simon's General Problem Solver (GPS) and Frank Rosenblatt's Mark I Perceptron hyped the market about AI potential and attracted many investments to the field. I believe to be important, however, to highlight that investments in AI research at the time were mostly deriving from public sources and carried out in universities. This led to huge results like the first program able to perform tasks as well as human would - the *geometry machine* developed by Nathaniel Rochester and Herbert Gelernter at IBM in 1959 - and the first AI chatbot ELIZA - developed by Joseph Weizenbaum at MIT in 1966.

The hype generated by such rapid advancements in AI quickly converted into a double-edged sward. Researchers were unable to meet the unrealistic promises set by the media and marketing agencies. In 1966 and 1973 respectively, the US and the UK governments - the two biggest funding sources for AI - drastically dropped their financial contributions to AI research. This set the beginning of the first AI winter. Approximately going from 1973 to 1980, this period was characterized by a slower, yet continuous, development for AI. Toosi et al., 2021 identify three main causes for this first AI winter:

1. the desire to merely replicate the way humans complete tasks instead of analysing and braking down the task itself into a usable algorithm (the so-called *thinking humanly* approach);

2. the underestimation of the complexity of certain tasks together with the non-linearity of scaling AI systems;

3. the negative public perception of AI that took ground in those years due to false promises previously hype by the media.

While the first two reasons identified above are, in fact, due to an industry-wide research approach (technological constraint), the third one falls in the sociocultural constraints presented in the previous section.

In the attempt to overcome the bias-variance trade-off, in the early 1980s, the industry shifted its focus to machines with narrower scopes in specific domains (known as *expert systems* approach). This nudged AI development into a new hype cycle attracting again huge investments. This time, however, private investors became more relevant than in the first summer as they began to see the direct impact that AI could have on their business. Alongside the progression in AI models, the 1980s saw a fast growth also in robotics, computer-vision system, and domain-specific hardware.

However, in between the end of the 1980s and the mid 1990s, a second AI winter arose. This time, still following Toosi et al., 2021, has to be found in the market dynamics. In fact, hardware manufactures failed to meet the ever-rising requests for specialized hardware made by the AI industry. This winter determined also quite a strong cultural shock making AI a cursed word for at least one decade. Research in AI, however, still did not come to a full stop. In fact, the second AI winter is well known for the backpropagation AI algorithm, the focus on statistics-based methods, and the creation of numerous public benchmark dataset still used nowadays.

In the early 1990s, the pace of AI development began to speed up again with wider applications in control theory, operational research, and statistics. Bayesian networks and Markov processes saw their first application in AI and reinforcement learning saw its first applications in robotics. But it was only when IBM's *Deep Blue* was able to defeat the chess world champion at a chess game in 1997 that AI affectively became mainstream again. At that time, hardware had improved quite a lot allowing to store and analyse more data from diverse sources and of multiple formats. Moreover, the computing power also increased following the well-known Moore's law and allowing to process all this data. The era of *Big Data* was just starting, and more machines confronted themselves in public directly with humans. In 2011, IBM's *Watson* won the TV show Jeopardy. The same year sees the tech-giant Apple launching its virtual assistant *Siri* bringing AI in the pockets of millions. In 2016, Google's *Alpha Go* won against the world's champion at the board game Go.

Another way to categorize AI's history up to these times, following Raximov et al., 2021, is to look at the mainstream trends in AI development.

1. The 1950s and 1960s — focused on the development of heuristic algorithms, have set the basis for more complex *General Problem Solvers*.

2. Between the late 1960s and the early 1970s — robotics took the lead with the development of sensory organs, knowledge, logical inference, behaviour strategy planning.

3. Since the late 1970s — AI research focused heavily on the creation of human-machine systems leveraging human intelligence and computer abilities.

This leads us more or less to the present where it seems like AI is living a new boom after the launch of *ChatGPT* by Open AI. The development of Generative Pretrained Transformers (GPTs) was boosted quite intensively attracting unprecedented levels of (mostly private) investments.

This brief history review of AI allowed us to develop important considerations on the automating technology under consideration, as well as its development cycles.
History shows that AI is quite complex and that its development requires a lot of highly diversified inputs to be simultaneously available. If only one of these inputs is scarce at a given point, AI development will eventually slow down into a new AI winter. At this point, the main question is: "how sustainable is the current rate of growth and adoption of AI technologies?".

This review highlights the fact that AI is in fact quite a broad term that has never reached a complete definition. Its meaning changed a lot over time and, in some periods, the term Artificial Intelligence has also been purposefully avoided by researchers and developers. It is important, thus, for the sake of this thesis, to provide the reader with a clear definition of the term AI (see Section 3.1.2).

3.1.2 A definition of AI

As simple and straightforward as it looks, the first definition of AI provided by McCarthy is itself incomplete and too generic as regards the objectives of this thesis. "The science and engineering of making intelligent machines" is an incomplete definition of AI since it does not specify the meaning of intelligence. When can a machine be intended to be intelligent?

In the attempt to reach a knowledge-based, coherent definition of AI we will follow the technical guidelines provided by the EU Joint Research Centre (JRC)[1], the US National Institute of Standards and Technologies (NIST)[2], the US Congressional Research Service[3], and the International Organization for Standardization (ISO)[4]. All experts are aligned on stating that AI is nowadays an umbrella term that requires a higher level of detail to be defined. There seems to be a general consensus in defining AI technologies with a white listing approach following a hierarchy similar to the one represented in the Venn diagram in Figure 3.1. Under this definition, Artificial Intelligence includes all the attempts at creating an artificial machine able to replicate human activities at a human-like level. These machines can either be highly specialized as in the case of the *expert systems* developed during the 1980s, or widely applicable to diverse tasks as *foundational models* or the most recent *agentic models*. The smallest set represents the ultimate goal for AI development: Artificial General Intelligence (AGI). This shows that Figure 3.1 can also be interpreted as a list of examples of AI models ordered from the least to the most complex and from the most to the least generalized. Every step down the ladder adds in complexity of the model whilst generalizing its applicability. This approach suggests that to define AI and its technical characteristics we need to analyse every step of this ladder one by one. A way to characterize AI models is to use three different metrics. Explainability represents how easy it is to understand the steps that a model takes in order to reach a solution. Traceability measures the ability to track back the history of data and decisions made by the model. Automatability is, instead, a proxy of which tasks can be automated by that technology. It must be noted that a higher level of explainability is usually associated with more traceability but less automatability. However, automatability and explainability are not fundamentally antithetical and there is active current research to make AI models more explainable.

- Expert Systems are computer programs that aim at replicating pre-

[1]Samoili et al., 2021.
[2]Theofanos et al., 2024.
[3]Harris and Wells, 2025.
[4]ISO, 2026.

viously observed human behaviour. At this stage programmers still code line by line every instruction that the machines is expected to carry out. This makes expert systems score highly on explainability and traceability while allowing them to automate only the specific tasks for which they have been programmed.

- Machine Learning (ML) opens up to a deeper level of generalization allowing the machine to learn from previous mistakes. ML models can be subdivided in three groups: supervised learning, unsupervised learning, and reinforcement learning.

  - Supervised learning is mostly used to subdivide data in pre-defined groups (classification models) or to forecast future trends of an existing time series (regression models). Common supervised learning algorithms are K-nearest-neighbours (KNN), Decision Trees, Random Forests, Support Vector Machines, Naive Bayes models, Gradient Boosting, and Regression models. All these models have different levels of explainability and traceability but a similar (low) level of automatability as they can only perform tasks for which there is enough pre-existing data on which to train the model on.

  - Unsupervised learning methods are commonly used to group data in organically defined groups (clustering models) or reducing the dimension of a dataset without any knowledge loss (data reduction models). Both applications are quite explainable and traceable as their results can easily be plotted in human-comprehensible diagrams, but the operator might still miss the reason why a specific cluster might be meaningful. As for automatability, unsupervised learning algorithms can cluster data without the need for previous knowledge, but their weights must be tuned ad-hoc for their specific use. These models are thus not highly generalized.

  - Reinforcement learning has wider application and represents algorithms that are able to interact with the environment they operate in and learn from the feedback received from such environment. Reinforcement learning algorithms usually score quite low in explainability and traceability but high in automatability. They suffer from the bias-variance trade-off during the learning period which can make them either more precise but less resilient to changes in the environment or less precise but more generalized.

- Deep Learning (DL) models are often referred to as Neural Networks due to their design inspired by the neural network of human brains. Compared to Machine Learning, Neural Networks are more complex and their behaviour is usually presented as a Black Box due to their low levels of explainability and traceability. They require huge amounts of data in order to work, but are able to generalize knowledge better than ML models as they are not dependent on predefined statistical algorithms.

- Generative Artificial Intelligence (Gen AI) identifies any AI model that is capable to generate any kind of content (text, audio, images, or videos). The most common type of Gen AI models are Generative Pretrained Transformers (GPTs). These models require significant amounts of training data from which to learn how to generate the requested data and, usually, are associated with a Natural Language Processing model to receive input in natural language from the user. They express high automatability potential for tasks that require a thought process as they mimic a human-like reasoning. However, they operate in a Black Box environment but can be trained to explain every step they take to reach a solution.

- Agentic models introduce the ability for AI of the models described above to act independently. Usually triggered by a user request, agentic models are able to perform action and interact with the digital or physical world. Their interaction with the physical world requires the interaction with Robotics. As of now, agentic models are still in development, but they promise automatability also for multi-process tasks that require reasoning and understanding of a changing environment. It is early, instead, to make some comments on explainability and traceability of these models but, most likely, they should be able to produce a track record of all the actions taken, and the data used to take such actions.

- Artificial General Intelligence (AGI) represents the end goal in AI development and promises an AI system that is autonomous and better than humans in any (old or new) task. Reaching AGI would present the ultimate level of automatability as it should allow for automation of any tasks imaginable, including technological development. As for what regards explainability and traceability, it is hard to make predictions on how it will behave. One concern is that AGI might reach a point of complexity called *singularity*. Technological singularity occurs when technology develops beyond human comprehension.

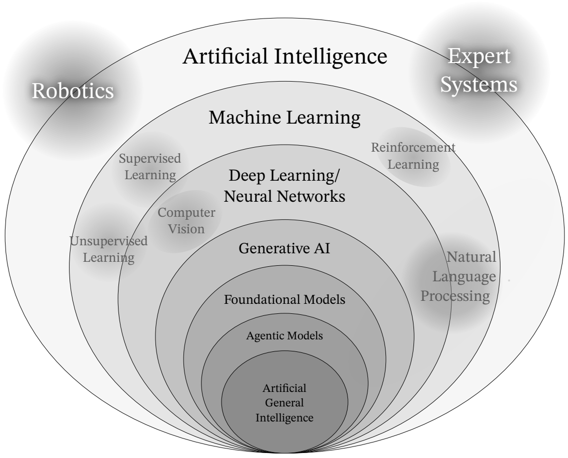

FIGURE 3.1: CLASSIFICATION OF AI TECHNOLOGIES

In conclusion, if the goal is to find a technical definition of AI, we must instead split the definition in the different applications for AI. If the goals, instead, is to a find a definition valid for economic research, literature suggests generalizing AI with one of three terms: automation, machines, robots. While automation has a more general meaning that can be expanded also to non-AI technologies, robots is usually used to refer to AI technologies able to interact with a physical environment. Instead, the expression 'machines' is somewhat in the middle of the two terms and is usually used to indicate any automation able to interact also indirectly with the physical world.

3.2 Comparing AI with other automation

The recency and the saliency of AI news might amplify the effects the technology truly has on our lives. This amplification regards both positive and negative effect AI has in the different markets that compose our economy making it difficult to evaluate whether different promises of AI's potentials are plausible or not. The previous historical and technical analysis of AI aimed at cancelling all the background noise generated by this amplification effect to objectively classify potential differences AI has with respect to other historical automation technologies.

It is true that, by definition, the goal of AI is to replicate and exceed human behaviour. A goal that, however, is quite similar to the idea of automation: building machines able to substitute humans in task production. What seems so different this time is that AI is somewhat redefining what we mean for *human behaviour*. With this I mean those behaviours that are typically and exclusively human. Until now a distinguishably human behaviour has been speech. We are the only animal species that uses speech (and text) to communicate. Now computers are able to reproduce this behaviour with a quality that is mostly undistinguishable from real human works. So, if it is not speech that makes us humans, what does? This is not an easy question, and rather than looking for an answer to it, we should wander why it is so relevant for us. Finding out what distinguished humans from machines will just give the incentive to innovate and build new machines that resemble humans more than the previous ones. This is just in the nature of us humans: we have a deep desire to think big. And what's bigger than being able to replicate the complex and perfect equilibrium of nature?

This is just to say that I strongly believe some kind of AGI will be achieved in future. However, it is hard for me to believe that achieving such level of automation will ever stop humanity from working, thinking, and creating. Even though AI might, at first, seam to bring a different type of automation in the labour market as it is replacing workers in tasks that we believed to be deeply human, in reality it represents just the natural cause of humans' ambition to think big. Moreover, we should take into consideration the four essential conditions for innovation listed at the end of chapter 1. An automation technology to be implemented requires being simultaneously technologically implementable, financially viable, environmentally acceptable, and socially recognized. If only one of these pillars falls, a full implementation of the technology is impossible. The first AI winter, for example, was caused by a lack of sociocultural support for the technology which, in turn, lowered investments. The second AI winter, instead, was mostly caused by the lack of technology which led companies to unfulfill their claims on what AI could do for their customers. It could be that the next AI winter will be caused by environmental and financial limitations. Even with the possibility of another AI winter, however, the shadow of AGI persists and economists are starting to develop some models able to hypothesize what an economy where work is not strictly required would operate[5,6].

[5] Restrepo, 2025.
[6] Bara, 2025.

Bibliography

Acemoglu, D., & Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of labor economics* (pp. 1043–1171, Vol. 4). Elsevier.

Acemoglu, D., & Restrepo, P. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American economic review*, *108*(6), 1488–1542.

Acemoglu, D., & Restrepo, P. (2022). Tasks, automation, and the rise in us wage inequality. *Econometrica*, *90*(5), 1973–2016.

Acemoglu, D., & Restrepo, P. (2024). *Automation and rent dissipation: Implications for wages, inequality, and productivity* (tech. rep.). National Bureau of Economic Research.

Aghion, P., Jones, B. F., & Jones, C. I. (2017). *Artificial intelligence and economic growth* (tech. rep.). National Bureau of Economic Research.

Autor, D. H. (2015). Why are there still so many jobs? the history and future of workplace automation. *Journal of economic perspectives*, *29*(3), 3–30. https://doi.org/10.1257/jep.29.3.3

Autor, D. H., Levy, F., & Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly journal of economics*, *118*(4), 1279–1333.

Bara, M. (2025). Moravec's paradox and restrepo's model: Limits of agi automation in growth. *arXiv preprint arXiv:2509.24466*.

Bowen, H. R. (1966). *Report of the national commission on technology, automation, and economic progress: Volume i*. Washington, U.S. Government Printing Office.

Chui, M., Manyika, J., & Miremadi, M. (2015). Four fundamentals of workplace automation. *McKinsey Quarterly*, *29*(3), 1–9.

Chui, M., Manyika, J., & Miremadi, M. (2016). Where machines could replace humans-and where they can't (yet). *The McKinsey Quarterly*, 1–12.

Costinot, A., & Vogel, J. (2009). Matching and inequality in the world economy. *Journal of Political Economy*, *118*(4), 747–786.

Duckworth, P., Graham, L., & Osborne, M. (2019). Inferring work task automatability from ai expert evidence. *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, 485–491. https://doi.org/10.1145/3306618.3314247

Ghasemi, Z., Afshar Kermani, M., & Allahviranloo, T. (2021). Exploring the main effect of e-banking on the banking industry concentration degree on predicting the future of the banking industry: A case study. *Advances in Fuzzy Systems*, *2021*(1), 8856990.

Harris, L., & Wells, N. (2025). Artificial intelligence (ai) taxonomy. (IG10077). https://www.everycrsreport.com/files/2025-04-25_IG10077_4b9ea4453ad7cc4f62e433aeb6f9c73ec443972b.pdf.

Hu, K. (2023). Chatgpt sets record for fastest-growing user base - analyst note. *Reuters*. https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/

ISO. (2026). *What is artificial intelligence (AI)?* Retrieved February 2, 2026, from https://www.iso.org/artificial-intelligence/what-is-ai

Keynes, J. M. (1930). Economic possibilities for our grandchildren. In *Essays in persuasion* (pp. 321–332). Springer.

Lu, Y., & Zhou, Y. (2021). A review on the economics of artificial intelligence. *Journal of Economic Surveys*, *35*(4), 1045–1072.

McCulloch, W. S., & Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. *The bulletin of mathematical biophysics*, *5*(4), 115–133.

OECD. (2025). OECD Compendium of Productivity Indicators 2025. https://doi.org/10.1787/b024d9e1-en

Raximov, N., Primqulov, O., & Daminova, B. (2021). Basic concepts and stages of research development on artificial intelligence. *2021 International Conference on Information Science and Communications Technologies (ICISCT)*, 1–4. https://doi.org/10.1109/ICISCT52966.2021.9670085

Restrepo, P. (2025). We won't be missed: Work and growth in the era of agi. *NBER Chapters*.

Rifkin, J., & Heilbroner, R. L. (1995). *The end of work*. Jeremy P. Tarcher.

Samoili, S., Lopez Cobo, M., Delipetrev, B., Martinez-Plumed, F., Gomez Gutierrez, E., & De Prato, G. (2021). Ai watch. defining artificial intelligence 2.0. (KJ-NA-30873-EN-N). https://doi.org/10.2760/019901

Theofanos, M., Choong, Y.-Y., & Jensen, T. (2024). Ai use taxonomy, a human-centered approach. https://doi.org/10.6028/NIST.AI.200-1

Toosi, A., Bottino, A. G., Saboury, B., Siegel, E., & Rahmim, A. (2021). A brief history of ai: How to prevent another winter (a critical review). *PET clinics*, *16*(4), 449–469.

Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, *59*(236), 33–60.

Turing, A. M. (1936). On computable numbers, with an application to the entscheidungsproblem. *J. of Math*, *58*(345-363), 5.

Zeira, J. (1998). Workers, machines, and economic growth. *The Quarterly Journal of Economics*, *113*(4), 1091–1117.