International Journal of Research and Review
[Indexed, Peer-Reviewed Journal]
Volume 13; Issue: 1; January 2026
Website: www.ijrrjournal.com
Review Paper
E-ISSN: 2349-9788; P-ISSN: 2454-2237

AI Agents and the Structural Transformation of Labour

Enrico Moch[1], Peter Schlecht[2]

[1]Academic Director, Department of Economics, GrandEdu Research School, Germany,
ORCID: https://orcid.org/0009-0005-4722-0961
[2]CEO, Department of Economics, GrandEdu GmbH, Germany,
ORCID: https://orcid.org/0009-0002-5371-750X

Corresponding Author: Enrico Moch, Email: Enrico.Moch@GrandEduResearchSchool.de

DOI: https://doi.org/10.52403/ijrr.20260132

ABSTRACT

This paper analyses the use of AI agents from a macroeconomic perspective. It shows that AI agents can trigger a structural break in the organisation of production and work. In contrast to earlier automation, AI agents not only increase the productivity of human labour. Rather, they are able to completely and permanently replace entire task areas. The central mechanism is a competition-driven pressure to automate. As soon as AI agents can fulfil tasks more cost-effectively than human labour, companies are forced to automate these tasks under competitive conditions in order to remain competitive. Automation is therefore not an optional innovation path, but an endogenous equilibrium outcome. The paper develops a task-based macroeconomic modelling framework in which AI agents are modelled as an independent factor of production that competes directly with labour for tasks. Decreasing agent costs and increasing agent productivity systematically shift an automation threshold over time and lead to the complete substitution of labour in growing task clusters. The analysis shows that rising output and falling labour demand can occur simultaneously and that productivity gains do not necessarily go hand in hand with stable employment. Furthermore, it becomes clear that the functional distribution of income is structurally shifting from labour income to agent and capital income, which creates potential risks for aggregate demand. The paper discusses these results in the context of the existing macroeconomic literature on technological change and shows that established models do not adequately capture the competition-driven compulsion for complete task substitution. The contribution lies in the explicit modelling of this mechanism and in the reassessment of the relationship between competition, productivity and employment in the age of AI agents.

Keywords: *Labour market, automation, AI agents, macroeconomics, productivity, technological change, competitive dynamics, value creation*

INTRODUCTION

The enormous progress in AI agents marks a turning point in the relationship between technology, competition and labour (Acemoglu & Restrepo, 2018). While earlier waves of automation mainly affected individual activities or clearly defined routines (Autor et al., 2003; Autor & Dorn, 2013), AI agents are enabling the complete replacement of entire task areas for the first time (Acemoglu & Autor, 2011; Zeira, 1998). This development does not represent a gradual continuation of technological change, but a structural break in the organisation of production and work (Aghion & Howitt, 1992). The key driver of this rupture is not just the technical performance of AI agents, but their economic cost structure (Acemoglu & Restrepo, 2018). AI agents are scalable, replicable and associated with very low marginal costs (Schmidhuber, 2015). As soon as they can fulfil a task more cost-effectively than human labour, the cost function of companies changes permanently (Grossman & Rossi-Hansberg, 2008; Zeira, 1998). Under competitive conditions, this change is imperative (Aghion & Howitt, 1992). Companies that forego automation produce at higher unit costs and lose market share (Grossman & Rossi-Hansberg, 2008). In this context, automation is not a strategic option, but an equilibrium result of competition (Romer, 1990; Aghion & Howitt, 1992). A simple example illustrates this mechanism. Two competing companies produce an identical good and perform a specific task using human labour. By using an AI agent, this task can now be automated at a lower cost (Acemoglu & Restrepo, 2018). If only one of the companies automates, its production costs fall (Grossman & Rossi-Hansberg, 2008). It can lower the price, gain market share and make profits (Romer, 1990). The other company has no real choice. It must also automate or exit the market (Aghion & Howitt, 1992). The automation of this task is thus inevitably spreading (Zeira, 1998). It is not the result of individual preferences, but the direct consequence of relative cost pressure under competition (Grossman & Rossi-Hansberg, 2008). This logic is macroeconomically relevant because it works across the board (Baqaee & Farhi, 2019). As soon as AI agents are more cost-effective for certain tasks, these tasks are fully automated in equilibrium (Zeira, 1998; Acemoglu & Restrepo, 2018). Labour is not supplemented or upgraded in these tasks, but displaced (Acemoglu & Restrepo, 2020). The use of AI agents therefore does not primarily lead to a reorganisation of work, but to its substitution (Zeira, 1998). The existing macroeconomic literature does not adequately capture this mechanism (Acemoglu & Autor, 2011). Classical models of technological change treat productivity progress as an exogenous shock or as a continuous process (Solow, 1956; Lucas, 1988). Approaches to skill-biased technological change or routine-based automation explain shifts within labour demand (Autor et al., 2003; Autor & Dorn, 2013), but implicitly assume that labour as a factor of production is fundamentally preserved (Acemoglu & Autor, 2011). AI agents call this assumption into question (Acemoglu & Restrepo, 2018). If tasks can be fully and permanently automated, labour is no longer complementary, but replaceable (Zeira, 1998). This paper develops a task-based macroeconomic modelling framework to analytically capture this structural break (Acemoglu & Autor, 2011; Grossman & Rossi-Hansberg, 2008). AI agents are modelled as an independent factor of production that competes with human labour for tasks (Zeira, 1998). The central mechanism is simple and robust. The most cost-effective technology is used for each task (Grossman & Rossi-Hansberg, 2008). Decreasing costs of AI agents shift an automation threshold so that more and more tasks are fully automated (Acemoglu & Restrepo, 2018). Under competition, this process is endogenous, accelerated and irreversible (Aghion & Howitt, 1992; Zeira, 1998). From a macroeconomic perspective, this has far-reaching consequences. Productivity gains can be accompanied by a decline in labour demand (Acemoglu & Restrepo, 2020; Graetz & Michaels, 2018). The distribution of value creation is systematically shifting from labour income to capital and agent rents (Fuest et al., 2019; Bofinger, 2019). At the same time, risks for overall economic demand arise when income gains are highly concentrated (Summers, 2014; Basu & Fernald, 2000). The use of AI agents is thus becoming a central determinant of macroeconomic dynamics and not a marginal technology issue (Baqaee & Farhi, 2019). The contribution of this paper lies in the explicit modelling of this competition-driven substitution mechanism. Automation is not understood as an exogenous shock or a slow adjustment process (Solow, 1956), but as a necessary equilibrium outcome under competition (Aghion & Howitt, 1992; Romer, 1990). The analysis shows that complete task substitution is not an extreme case, but the logical consequence of falling agent costs (Zeira, 1998). This mechanism is expected to operate most strongly in contestable markets, where unit-cost differences are rapidly translated into price pressure and market selection. In regulated, protected, or highly concentrated sectors, diffusion may be delayed or partially constrained, even if persistent cost advantages eventually shift adoption incentives. The paper is structured as follows. First, the existing macroeconomic literature on automation and technological change is systematically categorised and its limitations in the context of AI agents are identified (Acemoglu & Autor, 2011; Acemoglu & Restrepo, 2018). A task-based modelling framework is then developed that explicitly integrates AI agents (Grossman & Rossi-Hansberg, 2008; Zeira, 1998). Building on this, the macroeconomic effects of full task substitution are analysed (Acemoglu & Restrepo, 2020; Baqaee & Farhi, 2019) and examined in greater depth in a two-sector structure (Jones, 2011). Finally, the implications for employment, income distribution and economic governance are discussed (Summers, 2014; Truger, 2012).

STATE OF MACROECONOMIC RESEARCH ON AUTOMATION AND TECHNOLOGICAL CHANGE

The following chapter classifies the relevant macroeconomic model families that analyse technological substitution, automation and labour market effects (Solow, 1956; Lucas, 1988; Acemoglu & Autor, 2011; Acemoglu & Restrepo, 2018). The focus is on those theoretical building blocks that are central to AI agents: the substitutability of labour, the endogeneity of technology adoption and the role of competition (Romer, 1990; Aghion & Howitt, 1992; Grossman & Rossi-Hansberg, 2008).

The key question is which parts of the existing literature already capture the mechanism of complete task substitution under competitive pressure and where a theoretical gap remains.

Neoclassical growth theory and exogenous technical progress

The starting point is the Solow model, in which technological progress appears as an exogenous productivity parameter (Solow, 1956). In its standard form, technological progress increases labour. It increases the effectiveness of labour, but leaves labour as a necessary factor of production (Solow, 1956). The core mechanism is growth through capital accumulation and exogenous increases in productivity, not the structural replacement of labour. A closely related family are neoclassical growth models with alternative forms of technical progress, such as capital-increasing or factor-neutral. These variants can shift the wage share, returns and factor price ratios (Lucas, 1988). However, they continue to model technological change parametrically. Technology changes productivity levels, not the production structure in the sense that labour becomes completely dispensable for certain tasks (Solow, 1956; Lucas, 1988). For AI agents, the limits of these models are clear. They provide a foundation for growth and factor prices, but do not contain a mechanism that depicts complete task substitution as an equilibrium outcome under competition (Acemoglu & Autor, 2011). Labour can lose relative importance, but is not displaced from specific production steps (Acemoglu & Autor, 2011).

Endogenous growth theory and technology as a result of economic decisions

The endogenous growth literature shifts the analytical focus from exogenous productivity shocks to innovation and knowledge accumulation as endogenous economic processes (Romer, 1990). In Romer-type models, growth is driven by research and development that generate new ideas and intermediate goods, while in Lucas-type models human capital formation and learning constitute the central growth mechanism (Romer, 1990; Lucas, 1988). Schumpeterian quality ladder models complement this perspective by explaining growth as the outcome of innovation competition and creative destruction (Aghion & Howitt, 1992). These models are relevant for the present work because they treat technology development as an endogenous process (Romer, 1990; Aghion & Howitt, 1992). Nevertheless, the typical structure remains that technological progress increases productivity or creates new varieties, while labour remains a factor of production (Lucas, 1988; Romer, 1990). Labour displacement can occur via factor prices and sector shifts, but the complete replacement of entire task clusters is not a central, formally anchored mechanism (Acemoglu & Autor, 2011). The point of reference to AI agents is therefore twofold. First, these models provide an analytical language to think endogenous acceleration (Aghion & Howitt, 1992). Second, the concrete substitution logic at the task level usually remains underdetermined (Acemoglu & Autor, 2011). This is precisely where a task-based framework comes in (Grossman & Rossi-Hansberg, 2008; Zeira, 1998).

Skill-biased technological change and wage inequality

A central strand of the macroeconomic labour market literature is skill-biased technological change (Autor et al., 2003; Acemoglu & Autor, 2011). In this family of models, technology increases the productivity of skilled labour relative to less skilled labour (Autor et al., 2003). Classic contributions come from Autor, Katz and Krueger as well as from the broad empirical literature on increasing skill premiums (Autor et al., 2003; Acemoglu & Autor, 2011). At the theoretical level, a production function is often used in which skilled and less skilled labour have different elasticities of substitution with respect to capital or technology (Acemoglu & Autor, 2011). The strength of this literature lies in the explanation of wage dispersion and employment shifts (Autor & Dorn, 2013). However, its limitation in the context of AI agents is that technological impact is primarily modelled as a shift within labour demand (Acemoglu & Autor, 2011).

Labour remains central in the aggregate, the adjustment channel is qualification, reallocation and investment in education, not complete substitution (Autor et al., 2003; Autor & Dorn, 2013). This difference is crucial for the thesis pursued here. AI agents can not only shift the relative demand for skill groups, but also take over tasks themselves (Acemoglu & Restrepo, 2018; Zeira, 1998). This goes beyond skill bias. It is a structural substitution possibility that can occur independently of skill categories as soon as tasks tilt in terms of cost (Zeira, 1998; Acemoglu & Restrepo, 2018).

Task-based and routine-based automation

A significant advance over skill-based approaches is the shift from analysing occupations to tasks (Autor et al., 2003; Acemoglu & Autor, 2011). The routine-based automation literature, in particular the work of Autor, Levy and Murnane and its later operationalisation by Autor and Dorn (Autor et al., 2003; Autor & Dorn, 2013), is decisive here. The central idea is that routine tasks can be codified algorithmically and are therefore particularly easy to automate, while non-routine, especially interactive and contextual activities remain complementary (Autor et al., 2003). Theoretically, this perspective is formalised in task-based models, for example by Zeira as well as in later macroeconomic task models (Zeira, 1998; Grossman & Rossi-Hansberg, 2008). The common core is that production is understood as a bundle of heterogeneous tasks and that technology can substitute or complement tasks (Acemoglu & Autor, 2011; Grossman & Rossi-Hansberg, 2008). This literature is the natural starting point for

AI agents because it chooses the right unit of analysis (Acemoglu & Autor, 2011). At the same time, in many routine-based frameworks, substitution remains selective. The range of tasks that can be automated is limited or grows only slowly, and adoption is often modelled as gradual and partly discretionary (Autor et al., 2003; Autor & Dorn, 2013; Acemoglu & Autor, 2011). The logic of competition appears to be a background condition rather than a compelling equilibrium mechanism (Grossman & Rossi-Hansberg, 2008). AI agents fundamentally shift these parameters (Acemoglu & Restrepo, 2018). They expand the range of automatable tasks and significantly reduce marginal costs (Acemoglu & Restrepo, 2018; Zeira, 1998). As a result, the automation threshold is not only shifted, but can move quickly and broadly (Zeira, 1998). The task-based structure is suitable, but the mandatory competition mechanism and the possibility of complete task substitution must be modelled more explicitly (Acemoglu & Autor, 2011; Zeira, 1998).

Technology adoption, complementarities and capitalisation of automation

Another relevant strand of literature analyses technology adoption as an investment decision under adjustment costs, uncertainty and complementarities (Romer, 1990; Aghion & Howitt, 1992). This includes models in which new technologies only become productive after a reorganisation of production processes and in which switching costs delay adoption (Acemoglu & Autor, 2011). In the empirical literature, the observation that technology effects are strongly dependent on management practices, process design and data infrastructure also plays a role (Basu & Fernald, 2000). In this context, it is of central macroeconomic importance that technology adoption does not take place immediately and completely, even if a technology is technically available (Acemoglu & Autor, 2011). This can create transition dynamics that stabilise employment in the short term or delay sectoral adjustments (Baqaee & Farhi, 2019). For the analysis of AI agents, this point must be categorised precisely. Adjustment costs can postpone the timing of automation, but they do not remove the underlying forcing mechanism once the cost advantage is sufficiently large (Grossman & Rossi-Hansberg, 2008; Zeira, 1998). In this case, competitive pressure acts as a selection mechanism (Aghion & Howitt, 1992). Firms that do not adopt cost-reducing technologies lose competitiveness and exit the market (Grossman & Rossi-Hansberg, 2008). Persistent cost disadvantages are incompatible with competitive viability. Adjustment costs affect the timing of automation, not its long-run direction, which is determined by relative costs and market selection (Acemoglu & Autor, 2011; Zeira, 1998).

Robots, AI and the recent macroeconomic automation literature

The diffusion of robotics and AI has given rise to a macroeconomic literature that treats automation as a labour-replacing technology (Zeira, 1998; Acemoglu & Restrepo, 2018). Core contributions analyse task displacement and task creation. The empirical robotics literature documents the associated effects on employment and wages (Graetz & Michaels, 2018; Acemoglu & Restrepo, 2020). These studies are central to the present analysis because they take automation seriously as an independent mechanism and partly endogenise it (Acemoglu & Restrepo, 2018). Particularly important in this context is the idea that technological innovation both replaces tasks and creates new tasks (Acemoglu & Restrepo, 2018). Within this framework, employment can remain stable as long as the creation of new tasks sufficiently compensates for the substitution effect. However, the open question in the context of AI agents is whether task creation can systematically compensate for the substitution effect if AI agents take on a large and growing range of tasks at low cost (Zeira, 1998; Acemoglu & Restrepo, 2018). Many models allow for new tasks as a structural counterforce without justifying that they necessarily arise to the same extent in the AI agent context (Acemoglu & Autor, 2011). This paper therefore deliberately positions itself on the side of a substitution constraint and treats task creation not as an automatic stabiliser, but as a condition to be empirically tested (Zeira, 1998; Acemoglu & Restrepo, 2018).

Competition, market dynamics and the lack of a coercive mechanism

Although competition is almost always present across the model families mentioned, it is rarely formalised as the central cause of complete automation (Solow, 1956; Lucas, 1988; Romer, 1990). In many models, technology adoption appears as an optimisation decision involving costs, risks and adaptation (Acemoglu & Autor, 2011). In what follows, this coercive mechanism is therefore understood as a structural tendency under sufficiently strong competition, rather than as a universal claim for all market environments. Where entry barriers, regulation, switching costs, or market power weaken competitive discipline, the timing and extent of automation can differ, but the relative-cost logic remains the benchmark for the direction of long-run adjustment. Competition influences profits, but not necessarily the binary question of whether a task is automated in equilibrium as soon as relative cost ratios tilt (Grossman & Rossi-Hansberg, 2008). For AI agents, precisely this point is crucial (Acemoglu & Restrepo, 2018). If agent input for a task becomes cost-effective, non-adoption becomes a dominance trap (Aghion & Howitt, 1992; Zeira, 1998). The economic logic of selection generates widespread diffusion (Zeira, 1998). This creates a qualitatively different mechanism than in many traditional frameworks. Automation is not only possible, but necessary in order to survive in the market under competitive conditions (Aghion & Howitt, 1992; Grossman & Rossi-Hansberg, 2008).

Interim conclusion and theoretical gap

The existing literature provides three building blocks that are indispensable for analysing AI agents. First, it provides the macroeconomic language of growth, factor prices and income distribution (Solow, 1956; Lucas, 1988). Second, it provides a task-based perspective that correctly maps technological impact to tasks rather than occupations (Autor et al., 2003; Acemoglu & Autor, 2011; Grossman & Rossi-Hansberg, 2008). Thirdly, it shows that automation can directly reduce labour demand and that new tasks do not automatically compensate for this effect (Zeira, 1998; Acemoglu & Restrepo, 2018). At the same time, a central theoretical gap remains. There is a lack of an explicit macroeconomic framework that models AI agents as an independent production factor with low marginal costs and in which competition generates an endogenous automation constraint that leads to the complete replacement of entire task areas (Acemoglu & Autor, 2011; Zeira, 1998). The model framework developed in the next chapter addresses precisely this gap by formalising task selection via relative unit costs and deriving automation as an equilibrium outcome under competition (Grossman & Rossi-Hansberg, 2008; Aghion & Howitt, 1992).

CONCEPTUAL FRAMEWORK: KI AGENTS AS AN INDEPENDENT PRODUCTION FACTOR

This chapter develops the conceptual framework on which the formal modelling is based. The aim is to clearly locate AI agents in economic terms and to distinguish them clearly from existing production factors and technologies (Acemoglu & Autor, 2011). The focus is not on technical details, but on the characteristics of AI agents that are macroeconomically relevant (Schmidhuber, 2015).

Production factors and technology in macroeconomics

In traditional macroeconomics, production is described by the combination of a few basic factors, typically labour and capital (Solow, 1956). Technology appears either as a productivity parameter or as a process that increases the efficiency of these factors (Solow, 1956; Lucas, 1988). In both cases, the production structure remains essentially stable. Technology affects factors, but does not completely replace them (Solow, 1956; Lucas, 1988). This view is inadequate for AI agents (Acemoglu & Autor, 2011). AI agents are neither mere productivity enhancements of existing factors nor classical capital in the narrower sense (Zeira, 1998). They take over specific production functions that were previously performed by labour and can perform these functions completely without the need for complementary human inputs (Acemoglu & Restrepo, 2018). This creates the need to model AI agents as an independent production factor that competes directly with labour for tasks (Zeira, 1998).

Differentiating AI agents from capital and traditional automation

Traditional capital increases the productivity of labour, for example through machines, infrastructure or tools (Solow, 1956). Even with high capital intensity, labour remains necessary, be it to operate, monitor or control the production process (Lucas, 1988). Classic automation follows this logic. It replaces certain activities, but creates new complementary tasks (Autor et al., 2003; Autor & Dorn, 2013). AI agents differ from this in three key respects. AI agents exhibit three defining characteristics. They can perform tasks autonomously without continuous human control (Schmidhuber, 2015). They are highly scalable and replicable, allowing additional production units to be deployed at very low marginal cost (Schmidhuber, 2015; Zeira, 1998). They can also perform cognitive and coordinative tasks that were previously considered non-automatable (Acemoglu & Restrepo, 2018). These characteristics mean that AI agents not only complement labour, but can also completely replace it at the task level (Zeira, 1998). What is economically relevant here is not autonomy per se, but the fact that the marginal input of human labour can drop to zero for certain tasks (Zeira, 1998).

Tasks as a central unit of analysis

In order to map this substitution possibility, the conceptual framework shifts the analysis from occupations and activities to tasks (Autor et al., 2003; Acemoglu & Autor, 2011). Tasks are clearly defined production units that are necessary to generate output (Grossman & Rossi-Hansberg, 2008). In principle, each task can be fulfilled by different technologies (Zeira, 1998). This perspective allows technological substitution to be defined precisely (Acemoglu & Autor, 2011). Labour is not replaced in the abstract, but specifically where a task can be performed more cost-effectively by an AI agent (Zeira, 1998). The decision is made at task level and aggregates to macroeconomic effects (Baqaee & Farhi, 2019). The task approach is therefore not a mere modelling choice, but a necessary prerequisite for a clean analytical representation of complete substitution (Acemoglu & Autor, 2011; Zeira, 1998).

Cost structure and relative prices

The central economic mechanism results from the cost structure of AI agents (Zeira, 1998). While labour is limited by wages and its scaling can be associated with rising costs (Lucas, 1988), AI agents typically have falling average costs and very low marginal costs (Schmidhuber, 2015; Zeira, 1998). What is macroeconomically relevant here is not the absolute cost level, but the ratio of the costs of AI agents to the costs of human labour (Grossman & Rossi-Hansberg, 2008). As soon as an AI agent can perform a task more cheaply than human labour, the relative advantage is permanently reversed (Zeira, 1998). This cost relationship is the key to the automation threshold (Acemoglu & Restrepo, 2018). It determines which tasks are automated and which remain performed by human labour (Zeira, 1998). Falling agent costs or rising wages systematically shift this threshold over time, thereby redefining the cost-minimising task allocation in equilibrium (Acemoglu & Restrepo, 2018).

Competition as a compelling mechanism
Competition operates as a selection mechanism that enforces this allocation. Firms are not choosing whether to automate, but are selected according to whether they adopt the lowest-cost technology available for each task. Non-adoption implies persistently higher unit costs and is therefore incompatible with a competitive equilibrium. Automation thus emerges as a necessary equilibrium outcome rather than a discretionary firm-level decision (Aghion & Howitt, 1992; Grossman & Rossi-Hansberg, 2008).Under competitive conditions, every avoidable cost saving leads to price pressure (Romer, 1990). Companies that do not adopt lower-cost technologies lose market share or exit the market (Aghion & Howitt, 1992). Automation thus becomes an equilibrium outcome (Zeira, 1998). It is not dependent on preferences, norms or institutional arrangements, but on relative costs (Grossman & Rossi-Hansberg, 2008). Competition acts as a selection mechanism that enforces AI agents where they are more cost-effective than human labour (Aghion & Howitt, 1992; Zeira, 1998). This logic is decisive for the macroeconomic character of the analysis. Automation does not occur selectively, but across the board as soon as a critical cost ratio is reached (Zeira, 1998; Acemoglu & Restrepo, 2018).

From individual adoption to macroeconomic structural change
From a microeconomic perspective, the use of an AI agent is a cost minimisation decision (Grossman & Rossi-Hansberg, 2008). From a macroeconomic perspective, these decisions aggregate into a structural change in the mode of production (Baqaee & Farhi, 2019). Entire task areas are automated and the demand for labour does not fall gradually, but discretely in the affected task clusters (Zeira, 1998). The conceptual framework makes it clear that this change is not the result of political decisions or individual corporate strategies. It is the logical consequence of relative costs under competition (Aghion & Howitt, 1992; Grossman & Rossi-Hansberg, 2008). AI agent automation thus differs fundamentally from earlier forms of technological change, in which work was predominantly supplemented or reorganised (Autor et al., 2003; Acemoglu & Autor, 2011).

Transition to formal modelling
The conceptual considerations in this chapter provide the basis for the formal modelling in the next section. AI agents are introduced as an independent production factor; tasks are defined as the central unit of analysis and competition is anchored as a mandatory automation mechanism. The following chapter translates this framework into a task-based macroeconomic model in which the complete replacement of tasks by AI agents is derived as an equilibrium outcome (Zeira, 1998; Acemoglu & Autor, 2011).

FORMAL MODELLING: TASK-BASED PRODUCTION WITH KI AGENTS

This chapter develops the formal macroeconomic model that explicitly depicts the competition-driven automation constraint through AI agents. The aim is to show under which conditions tasks are completely replaced by AI agents and how this process aggregates to macroeconomic variables (Zeira, 1998; Grossman & Rossi-Hansberg, 2008). The model deliberately focuses on a limit case in which tasks are allocated to the lowest-unit-cost technology, in order to make long-run structural implications transparent. Hybrid arrangements and transitional forms of human-agent collaboration are important in practice, but are treated here as adjustment paths around the core selection mechanism, not as substitutes for it.

Basic structure of the economy
The economy consists of a large number of identical firms, a representative household and competitive factor markets. Time is initially static. There are two factors of

production: human labour and AI agents. AI agents are modelled as marketable inputs that can be purchased at a given price. Companies produce a homogeneous end product.
They are price takers and minimise production costs for a given level of output (Grossman & Rossi-Hansberg, 2008). Competition ensures that more cost-effective production methods are used across the board in equilibrium (Aghion & Howitt, 1992).

Task-based production technology
Production is modelled as an aggregation of a continuum of tasks. Tasks are indexed by i ∈ [0,1]. Each task is necessary to produce output. The aggregated output Y results from a CES aggregation of the completed tasks:

$$Y = \left(\int_0^1 y(i)^{\frac{\sigma-1}{\sigma}} di \right)^{\frac{\sigma}{\sigma-1}}, \sigma > 1$$

The parameter σ determines the elasticity of substitution between tasks. This specification allows a clean aggregation of tasks into output and is standard in task-based macroeconomics (Dixit & Stiglitz, 1977; Grossman & Rossi-Hansberg, 2008; Acemoglu & Autor, 2011).

Technologies at task level
There are two alternative technologies for each task i. The task can be fulfilled either by human labour or by the use of AI agents. The effective task fulfilment results from the technology that delivers the highest effective task performance and is defined as:

$$y(i) = \max\{ A_{L(i)} * l(i), A_A(i) * a(i)\}$$

Here, $l(i)$ denotes the use of labour on task i, and $a(i)$ the use of agent input. The parameters $A_L(i)$ and $A_A(i)$ are task-specific productivities of the respective technology. This specification implies complete substitutability at task level (Zeira, 1998; Acemoglu & Autor, 2011). For each task, only the technology that delivers the highest effective task performance is employed (Grossman & Rossi-Hansberg, 2008). If a task is automated, the optimal use of labour on that task is zero. Labour and AI agents are therefore not complementary, but stand in a direct substitution relationship (Zeira, 1998). The decision between labour and AI agents is made at the task level and aggregates to macroeconomic effects across all tasks (Acemoglu & Autor, 2011; Baqaee & Farhi, 2019).

Cost structure and automation conditions
The wage rate is ω. The price of a unit of agent input is p_A. The unit costs of task fulfilment are given by the ratio of input prices to task-specific productivities.

$$c_{L(i)} = \frac{\omega}{A_{L(i)}},$$

The unit costs for the use of AI agents are as follows

$$c_{A(i)} = \frac{p_A}{A_{A(i)}}$$

Companies choose the technology that minimises costs for each task. A task is automated if the unit costs of using the agent are not higher than those of human labour:

$$C_A(i) \leq C_L(i)$$

This condition can be written equivalently as

$$\frac{p_A}{A_A(i)} \leq \frac{\omega}{A_L(i)}$$

This inequality represents the central automation condition of the model. It shows that automation depends exclusively on relative costs and task-specific productivities. Technological progress in AI agents has an effect via $A_A(i)$, wage developments via ω, and price changes in agent input via p_A.

Organisation of tasks and automation threshold
To obtain a clear automation threshold, we define the relative productivity advantage of AI agents at the task level as

$$\theta(i) = \frac{A_{A(i)}}{A_{L(i)}}$$

We assume that $\theta(i)$ is not decreasing in i. Tasks with a low index i are therefore relatively easier to automate than tasks with a high index.
Under this assumption, there is a clear automation threshold $i^* \in [0, 1]$, such that:
$i \in [0, i^*]$ are automated, $i \in (i^*, 1]$ are performed by human labour.
The threshold i^* is implicitly defined by the indifference condition between human labour and agent deployment:

$$\frac{p_A}{A_A(i^*)} = \frac{\omega}{A_L(i^*)}$$

If the relative price of agent input p_A/ω falls, the automation threshold i^* rises and the automated task set expands accordingly. This monotonic relationship forms the core of the automation mechanism in the model.

Cost minimisation and marginal costs
The minimum marginal costs of production result from a task-level marginal cost index. The following applies to the CES structure:

$$c(i) = \min\{c_L(i), c_A(i)\}$$

Under perfect competition, the price of goods equals marginal cost (Solow, 1956). Under monopolistic competition, a markup is added that raises the price above marginal cost (Romer, 1990). In both cases, the choice of task remains identical, as it is determined by cost minimisation at the task level and not by pricing at the goods level (Grossman & Rossi-Hansberg, 2008). Competition here does not primarily work via prices, but via cost minimisation. As soon as the use of AI agents for a task is more cost-effective than human labour, this task is automated regardless of the market form (Zeira, 1998; Aghion & Howitt, 1992).

Aggregated factor demand
Labour demand arises exclusively in the non-automated task range $(i^*, 1]$(Zeira, 1998; Acemoglu & Autor, 2011). Aggregate labour demand is therefore a monotonically decreasing function of the automation threshold i^*. If the price of agent input p_A falls or the productivity of AI agents $A_A(i)$ increases, the automation threshold i^* shifts upwards (Zeira, 1998; Acemoglu & Restrepo, 2018). As a result, the range of tasks performed by human labour shrinks, and labour demand declines discretely as tasks become automated.
The demand for agent input is the mirror image of the automated task interval $[0, i^*]$. As automation expands, agent demand increases correspondingly. This structure makes clear that rising aggregate output and falling labour demand can occur simultaneously.

Competition and endogenous automation constraints
Under competitive conditions, automation is not an optional decision (Aghion & Howitt, 1992). For each task i, or which $c_A(i) < c_L(i)$ the use of AI agents minimises costs and is therefore imperative (Grossman & Rossi-Hansberg, 2008; Zeira, 1998). Companies that do not use this technology produce at higher marginal costs and are not viable in a competitive equilibrium (Aghion & Howitt, 1992; Grossman & Rossi-Hansberg, 2008). The complete automation of all tasks that can be automated in terms of costs is therefore not a strategic outcome, but an equilibrium outcome under competition (Zeira, 1998; Romer, 1990). This process is independent of expectations, norms or institutional preferences. It is determined solely by relative costs and competition.

Interim conclusion
The model shows that AI agents lead to the complete replacement of entire task areas with a suitable cost structure. The mechanism is simple, robust and macroeconomically compelling. Automation arises endogenously from competition and relative prices. In the next chapter, this basic model is extended to explicitly analyse

distributional effects, demand effects and sectoral structural change.

FORMAL MODELLING: TASK-BASED PRODUCTION WITH KI AGENTS

The previous modelling describes a static equilibrium. However, it is crucial for the macroeconomic relevance of AI agents that their cost and performance structure is not constant, but dynamic. This chapter adds a time dimension to the basic model and shows how an endogenous automation process arises.

Time structure and dynamic perspective

Time is discrete and is indexed by $t = 0, 1, 2, \ldots.$ In each period, companies make cost minimisation decisions based on the current technology parameters (Grossman & Rossi-Hansberg, 2008). The basic production structure remains unchanged, but the costs and productivity of AI agents change over time. The central assumption is that AI agents are subject to learning, scaling and diffusion effects that reduce their effective costs or increase their productivity (Romer, 1990; Zeira, 1998; Acemoglu & Restrepo, 2018).

Dynamics of agent costs

The unit price for agent input $p_A(t)$ is time-dependent. It reflects computational costs, model availability, data access and organisational integration. We assume that the price for agent input falls over time:

$$p_A(t+1) < p_A(t)$$

This assumption is economically plausible. As AI agents become more widespread, economies of scale, learning curves and increased competition arise in the agent market (Romer, 1990; Aghion & Howitt, 1992). The marginal costs of additional agent input units fall, while fixed costs are spread over ever larger production volumes (Grossman & Rossi-Hansberg, 2008). The observed cost reduction is therefore not exogenous, but a direct consequence of the utilisation and diffusion of AI agents themselves (Zeira, 1998; Acemoglu & Restrepo, 2018).

Productivity dynamics at task level

We assume weakly increasing task-specific productivity of AI agents over time:

$$A_A(i, t+1) \geq A_A(i, t)$$

This improvement can result from data feedback, model improvements or increasing specialisation (Romer, 1990; Schmidhuber, 2015). Crucially, productivity gains do not have to be evenly distributed across all tasks (Acemoglu & Autor, 2011). As a rule, AI agents improve first in tasks that are already automated (Zeira, 1998; Acemoglu & Restrepo, 2018). This creates a reinforcement mechanism. Automated tasks provide data and experience that enable further productivity gains and open up additional tasks for automation (Romer, 1990; Acemoglu & Restrepo, 2018).

Dynamic automation threshold

The automation threshold is now time-dependent and results from the indifference condition at task level. The automation threshold $i^*(t)$ is implicitly defined by

$$\theta(i*(t), t) = \frac{pA(t)}{\omega(t)}$$

If the price of the agent input $pA(t)$ falls or the productivity of the AI agents $A_A(i, t)$, increases, the automation threshold $i^*(t)$ shifts upwards. The automated task area thus grows over time (Zeira, 1998; Acemoglu & Restrepo, 2018). This shift is monotonic as long as $p_A(t)$ fällt oder $A_A(i, t)$ continues to rise. Once tasks are automated, they are not re-humanised in the model. The process is therefore irreversible in the absence of exogenous increases in agent input costs or negative technology shocks (Zeira, 1998; Aghion & Howitt, 1992).

Endogenous acceleration of the automation process

The combination of falling agent costs and rising agent productivity generates an endogenous acceleration of the automation process. The more tasks are automated, the stronger the effects of learning and economies of scale. These effects further reduce the price of the agent input $p_A(t)$ or increase the productivity of the AI agents $A_A(i,t)$, whereby additional tasks are automated. Formally, this means

$$i^*(t+1) \geq i^*(t)$$

with strict inequality as long as technological progress takes place:

$$i^*(t+1) > i^*(t)$$

The automation process is therefore self-reinforcing (Aghion & Howitt, 1992; Zeira, 1998). It differs fundamentally from earlier automation technologies, in which automation typically represented a one-off or limited adaptation step (Solow, 1956; Autor et al., 2003). In the present model, automation is not a static transition, but a dynamic, endogenously accelerated process driven by falling agent costs, learning and economies of scale as well as competitive pressure (Romer, 1990; Acemoglu & Restrepo, 2018).

Dynamic demand for labour

The demand for labour is now explicitly time-dependent. In each period, labour is demanded exclusively in the non-automated task area $(i^*(t), 1]$ (Zeira, 1998; Acemoglu & Autor, 2011). As the automation threshold $i^*(t)$ increases, this task area systematically shrinks. Even if the aggregate output $Y(t)$ increases, the demand for labour can fall (Acemoglu & Restrepo, 2020; Graetz & Michaels, 2018). Productivity gains do not necessarily compensate for the decline in labour, as the substitution effect of automation can dominate (Zeira, 1998). This results in a central macroeconomic finding: economic growth and employment can be permanently decoupled (Solow, 1956; Acemoglu & Restrepo, 2020).

Dynamic competitive logic

Even in a dynamic context, competition remains the central driver of the automation process (Aghion & Howitt, 1992). In every period, companies that do not adopt available cost reductions have higher marginal costs than their competitors (Grossman & Rossi-Hansberg, 2008). This effect increases over time, as early cost advantages have a cumulative effect (Aghion & Howitt, 1992; Zeira, 1998). Companies that automate early benefit from lower production costs and growing market shares (Romer, 1990; Grossman & Rossi-Hansberg, 2008). Late automatisers come under increasing competitive pressure and lose market position (Zeira, 1998). The diffusion of AI agents therefore does not occur evenly, but in waves, triggered by the crossing of cost-side thresholds (Acemoglu & Restrepo, 2018; Zeira, 1998). The dynamics do not arise through coordination or planning, but through decentralised market selection (Aghion & Howitt, 1992).

Summary of the dynamic results

The dynamic extension of the model provides three key results. First, automation is a self-reinforcing process as long as agent costs fall or agent productivities rise (Aghion & Howitt, 1992; Zeira, 1998; Acemoglu & Restrepo, 2018). Secondly, competition leads to automation becoming irreversible as soon as a task becomes cost-prohibitive (Zeira, 1998; Grossman & Rossi-Hansberg, 2008). Thirdly, rising productivity and falling labour demand can occur together in the long term (Acemoglu & Restrepo, 2020; Graetz & Michaels, 2018). The dynamic model thus provides the theoretical basis for analysing the macroeconomic consequences of complete task substitution (Zeira, 1998; Acemoglu & Autor, 2011).

Transition to the results

The following chapter uses this model framework to systematically analyse the macroeconomic effects of dynamism. The focus is on productivity, employment, income distribution and aggregate demand under the conditions of an accelerated automation process.

ANALYSIS AND RESULTS

This chapter analyses the macroeconomic consequences of the model developed. The starting point is the mechanisms derived in chapter 4. The aim is to present the key results transparently and categorise their economic significance.

Automation and labour demand

The central result of the model is the systematic displacement of labour by AI agents at the task level (Zeira, 1998; Acemoglu & Autor, 2011). For each task, labour is demanded precisely when its costs are below the costs of using the agent (Grossman & Rossi-Hansberg, 2008; Zeira, 1998). Decreasing agent costs or increasing agent productivity shift the automation threshold upwards over time (Acemoglu & Restrepo, 2018). In macroeconomic terms, this means that the demand for labour does not shrink continuously, but discretely (Zeira, 1998). Task-level automation implies a corner solution for labour input. When a task is automated, labour input is optimally set to zero (Zeira, 1998). As automation expands across tasks, aggregate labour demand can fall even in the presence of output growth, breaking the conventional link between employment and production (Acemoglu & Restrepo, 2020; Graetz & Michaels, 2018; Basu & Fernald, 2000). This mechanism stands in contrast to standard growth models, where technological progress raises effective labour productivity without eliminating labour from the production process (Solow, 1956; Lucas, 1988).

Productivity and output dynamics

At the same time, overall economic productivity is increasing (Solow, 1956; Basu & Fernald, 2000). AI agents enable more efficient fulfilment of automated tasks (Acemoglu & Restrepo, 2018). Economies of scale and learning processes further reduce effective production costs (Romer, 1990; Acemoglu & Restrepo, 2018). The model thus shows a central macroeconomic finding. Rising output and falling employment are not an exception, but a structural result (Zeira, 1998; Acemoglu & Restrepo, 2020). Productivity gains do not result from more intensive labour input, but from the substitution of labour by AI agents (Zeira, 1998). This mechanism explains why classical relationships between employment and growth can become unstable in the context of AI agents (Solow, 1956; Lucas, 1988; Acemoglu & Restrepo, 2020).

Distribution effects and factor income

The displacement of labour has a direct impact on income distribution (Zeira, 1998; Acemoglu & Restrepo, 2020). Labour income decreases as fewer working hours are demanded (Acemoglu & Restrepo, 2020). At the same time, income from agent capital and complementary capital increases (Fuest et al., 2019). In the model, this shift is not a temporary phenomenon (Zeira, 1998). As long as automation progresses, the share of value added systematically shifts from labour to agent and capital income (Zeira, 1998; Fuest et al., 2019). The functional distribution of income thus changes permanently (Zeira, 1998). What is macroeconomically relevant is that this shift is not compensated for by wage adjustments (Acemoglu & Autor, 2011). Falling wages cannot stop substitution as long as agent costs remain relatively lower (Grossman & Rossi-Hansberg, 2008; Zeira, 1998).

Aggregate demand

The change in income has direct consequences for aggregate demand (Basu & Fernald, 2000; Summers, 2014). Labour income generally has a higher propensity to consume than capital and pension income (Bofinger, 2019; Fuest et al., 2019). If the labour share falls, aggregate demand can fall short of production potential (Summers, 2014). The model makes it clear that productivity gains alone are not a sufficient condition for demand growth (Basu & Fernald, 2000; Summers, 2014). Without corresponding redistribution or demand channels, automation can lead to weak demand, even with increasing efficiency (Summers, 2014; Truger, 2012). This creates a macroeconomic tension between supply dynamics and demand development (Summers, 2014; Baqaee & Farhi, 2019).

Structural change and sectoral effects

The two-sector expansion shows that workers can move from automated tasks to less automatable areas (Autor & Dorn, 2013; Acemoglu & Autor, 2011). However, this reallocation mechanism is limited (Autor & Dorn, 2013). It depends on the demand for services that cannot be automated, on institutional framework conditions and on sectoral differences in productivity (Basu & Fernald, 2000; Truger, 2012). The model implies two possible regimes. In a reallocation regime, employment can remain relatively stable while the economic structure shifts significantly (Autor & Dorn, 2013). In a displacement regime, automation predominates and employment declines permanently (Zeira, 1998; Acemoglu & Restrepo, 2020). Which dynamic dominates is not a purely technological question, but a macroeconomic one (Summers, 2014; Baqaee & Farhi, 2019).

Dynamic amplification and irreversibility

The dynamic extension identifies automation as a self-reinforcing structural process. Initial task automation generates learning effects and economies of scale that systematically lower agent costs and raise productivity, thereby shifting the automation threshold over time (Aghion & Howitt, 1992; Zeira, 1998; Romer, 1990; Acemoglu & Restrepo, 2018). Competition functions as the decisive transmission mechanism in this dynamic. It ensures that relative cost advantages do not remain confined to individual firms, but translate into economy-wide diffusion, preventing selective adoption and stabilising automation as a general equilibrium outcome (Aghion & Howitt, 1992; Grossman & Rossi-Hansberg, 2008). Once tasks have been automated, they are not taken over again by humans in the model (Zeira, 1998). The process is irreversible as long as no exogenous shocks fundamentally change the cost ratios (Baqaee & Farhi, 2019). This irreversibility distinguishes AI agents from earlier technologies with cyclical or reversible labour demand (Solow, 1956; Lucas, 1988).

Summary of the key findings

The analysis provides four key findings. Firstly, the use of AI agents under competitive conditions leads to the complete substitution of entire task areas (Zeira, 1998; Acemoglu & Restrepo, 2018). Secondly, productivity and output can increase while employment decreases (Acemoglu & Restrepo, 2020; Graetz & Michaels, 2018). Thirdly, the distribution of income shifts systematically from labour to agent and capital income (Zeira, 1998; Fuest et al., 2019). Fourth, risks arise for aggregate demand and macroeconomic stability (Summers, 2014; Basu & Fernald, 2000). These results form the basis for the subsequent discussion of normative and economic policy implications (Truger, 2012; Snower, 1993).

DISCUSSION

The results of this study suggest that the use of AI agents should not be interpreted as an incremental continuation of existing automation processes, but rather as a qualitatively new mechanism of macroeconomic change (Acemoglu & Restrepo, 2018). The central finding is that automation under the conditions modelled here is neither gradual nor selective, but competition-driven, comprehensive and structurally imperative (Zeira, 1998; Aghion & Howitt, 1992). This characteristic distinguishes AI agents fundamentally from earlier technologies and calls into question central assumptions of macroeconomic theory (Solow, 1956; Lucas, 1988). A key feature of the existing literature is the implicit expectation that technological progress will either be complementary to labour in the long term or will be compensated by the emergence of new jobs (Autor et al., 2003; Acemoglu & Autor, 2011). This expectation is historically plausible, but theoretically not compelling (Zeira, 1998). It is based on the characteristics of earlier technologies that made labour more productive without completely replacing it on a broad task basis (Autor & Dorn, 2013). The analysis of this paper shows that AI agents cross precisely this boundary (Acemoglu & Restrepo, 2018). They can perform tasks autonomously, scalably and at very low marginal costs (Schmidhuber, 2015; Zeira, 1998). Labour is not supplemented or transformed in these tasks, but becomes superfluous (Zeira, 1998). A central point of discussion is the role of competition. In many macroeconomic models, competition functions as a framework condition, not as an active driver of technological substitution (Solow, 1956; Romer, 1990). The results of this paper suggest that competition should be understood as a selection mechanism that forces automation as soon as relative costs tilt (Aghion & Howitt, 1992; Grossman & Rossi-Hansberg, 2008). Companies are not faced with a real choice (Zeira, 1998). Forgoing lower-cost AI agents leads to structural competitive disadvantages and is not sustainable in equilibrium (Grossman & Rossi-Hansberg, 2008). Automation is therefore not an expression of entrepreneurial preferences, but a result of market selection (Aghion & Howitt, 1992). This perspective has far-reaching implications for the interpretation of technology adoption. While many models describe technological diffusion as a gradual process with adaptation costs and uncertainty (Acemoglu & Autor, 2011), the framework developed here shows that AI agent automation emerges endogenously from cost and competitive dynamics (Acemoglu & Restrepo, 2018; Zeira, 1998). Adaptation costs can delay the process, but not cancel it out (Acemoglu & Autor, 2011). As soon as the cost advantage is sufficiently large, automation becomes widespread (Zeira, 1998). The process is therefore not only endogenous, but potentially accelerated and irreversible (Aghion & Howitt, 1992). Closely linked to this is the discussion about new tasks and activities. Some of the literature argues that technological substitution is regularly compensated for by the emergence of new tasks (Acemoglu & Restrepo, 2018). This study does not fundamentally question this compensation logic, but clearly relativises it (Zeira, 1998). New tasks do not automatically arise to the same extent as automated tasks disappear (Acemoglu & Autor, 2011). Hence, task creation is a theoretical possibility but an empirical question; it cannot be assumed to be sufficiently large or durably labour-intensive in the AI-agent context. Particularly in the context of AI agents, there is no guarantee that new activities are permanently labour-intensive or cannot be automated themselves (Acemoglu & Restrepo, 2018). Task creation is therefore not a structural counter-mechanism, but an empirical possibility with an uncertain extent (Zeira, 1998). The analysis also sheds new light on the relationship between productivity and employment. In traditional macroeconomic models, productivity gains go hand in hand with rising prosperity and stable employment in the long term (Solow, 1956; Lucas, 1988). In the framework developed here, this link is broken (Acemoglu & Restrepo, 2020). Productivity can increase while labour demand decreases (Graetz & Michaels, 2018; Acemoglu & Restrepo, 2020). This decoupling is not a transitional phenomenon, but a structural characteristic of AI agent automation (Zeira, 1998). Distribution and demand effects thus gain central macroeconomic significance (Summers, 2014). The functional distribution of income is particularly relevant in this context. The systematic shift from value creation to agent and capital income is not a side effect in the model, but a direct consequence of the substitution of labour (Zeira, 1998; Fuest et al., 2019). This shift directly influences consumption, investment and macroeconomic stability (Basu & Fernald, 2000; Summers, 2014). If productivity gains are not translated into broad income, demand deficits can arise that limit growth potential (Summers, 2014). These theoretical insights have direct relevance for economic practice. For companies, the results imply that the use of AI agents is not a long-term option, but a short-term competitive necessity as soon as cost ratios tilt (Grossman & Rossi-Hansberg, 2008; Aghion & Howitt, 1992). Companies that delay automation risk permanent competitive disadvantages (Zeira, 1998). Strategic decisions are thus shifting from the question of whether to automate to the question of how quickly and to what extent automation must be implemented (Acemoglu & Restrepo, 2018). For labour market actors and social partners, the results suggest that declines in employment should not be interpreted primarily as a skills problem (Autor et al., 2003; Autor & Dorn, 2013). They are an expression of a structural substitution process (Zeira, 1998). Retraining and further training remain important, but cannot systematically prevent complete task substitution (Acemoglu & Autor, 2011). The expectation that qualification policy alone will stabilise employment appears increasingly unrealistic under the conditions of AI agents (Acemoglu & Restrepo, 2020). There are also fundamental implications for economic policy. If automation is competition-driven and irreversible, traditional labour market policy instruments will reach their limits (Truger, 2012). Macroeconomic stabilisation shifts more strongly to questions of income distribution, demand protection and the institutional design of transfer systems (Summers, 2014; Bofinger, 2019). Technology policy is thus inextricably linked to distribution and stabilisation policy (Snower, 1993). Finally, the analysis also concerns regulatory and competition policy. The regulation of AI agents is not just a question of technological safety or ethical guidelines (Schmidhuber, 2015). It directly influences market structures, employment and income distribution (Bofinger, 2019). Regulation can delay or channel the automation process, but cannot prevent it completely as long as the underlying cost relationships remain (Zeira, 1998; Grossman & Rossi-Hansberg, 2008). To summarise, the discussion shows that AI agents require a reassessment of central macroeconomic relationships (Acemoglu & Autor, 2011). Labour, competition and technological progress have a different relationship than in previous models (Solow, 1956; Lucas, 1988). This paper contributes to this reassessment by modelling automation as a competition-driven structural break and systematically working out its macroeconomic consequences (Zeira, 1998; Acemoglu & Restrepo, 2018).

LIMITATIONS

This paper deliberately pursues a strongly abstracting macroeconomic approach (Solow, 1956; Lucas, 1988). This abstraction is necessary in order to clearly visualise the central mechanism of a competition-driven automation constraint (Aghion & Howitt, 1992; Grossman & Rossi-Hansberg, 2008). At the same time, this results in limitations that are essential for categorising the results. A first limitation concerns the degree of model abstraction. The formal analysis reduces the complexity of real production processes to tasks that are clearly performed either by labour or by AI agents (Autor et al., 2003; Acemoglu & Autor, 2011). This binary structure makes it possible to derive complete task substitution in an analytically clean way (Zeira, 1998), but ignores hybrid forms of human-machine interaction (Autor et al., 2003). In practice, transitional forms exist in which tasks are partially automated or performed by combinations of human and machine input (Autor & Dorn, 2013). The model thus captures the long-term limit case, not necessarily the short-term adaptation path (Acemoglu & Autor, 2011). Closely linked to this is the assumption of homogeneous competitive conditions. The model assumes effective competition in which cost advantages are quickly translated into market share (Aghion & Howitt, 1992). In real markets, however, there is market power, regulatory protection mechanisms, institutional inertia and coordination problems that can delay or fragment automation (Bofinger, 2019; Truger, 2012). These factors influence the pace and course of the automation process, but do not necessarily change its long-term direction (Zeira, 1998; Grossman & Rossi-Hansberg, 2008). The results of the work are therefore to be understood as structural statements, not as forecasts of short-term employment trends (Solow, 1956). A further limitation concerns the treatment of technological development. The dynamic component models falling agent costs and rising productivity as monotonic processes (Acemoglu & Restrepo, 2018). This assumption reflects long-term learning and economies of scale (Schmidhuber, 2015), but abstracts from possible technological dead ends, regulatory interventions or resource restrictions (Summers, 2014). Exogenous shocks, such as rising energy prices, geopolitical fragmentation or regulatory intervention, could change the cost ratios temporarily or permanently (Baqaee & Farhi, 2019). Such scenarios are not explicitly modelled in the model. The role of new tasks and activities is also deliberately modelled with restraint. The work does not treat task creation as an automatic equalisation mechanism (Zeira, 1998; Acemoglu & Restrepo, 2018). At the same time, it does not formally analyse the conditions under which new tasks could arise to a sufficient extent to compensate for substitution effects (Acemoglu & Autor, 2011). Explicitly endogenising new tasks would extend the model framework, but could compromise the analytical clarity of the central mechanism (Acemoglu & Autor, 2011). The present work therefore prioritises analytical transparency over maximum completeness (Lucas, 1988). Another limitation concerns the demand and distribution perspective. Although income shifts and demand risks are discussed (Fuest et al., 2019; Summers, 2014), they are not modelled in a fully closed macroeconomic equilibrium with explicit monetary, fiscal and financial market integration (Basu & Fernald, 2000). In particular, feedbacks between automation, asset prices and financial stability remain outside the formal framework (Summers, 2014). These aspects are relevant for a comprehensive macroeconomic assessment, but would significantly shift the focus of the work (Schularick & Taylor, 2012). Finally, the work is deliberately theory-led and not empirically calibrated. The model makes no claim to short-term predictive accuracy (Solow, 1956). Its contribution lies in the identification of a mechanism, not in the quantitative estimation of its strength (Zeira, 1998). Empirical validation, sectoral differentiation and country-specific analyses are therefore reserved for future research (Acemoglu & Restrepo, 2020). The limitations of the work are to be understood as deliberately drawing boundaries.

They serve to clearly identify a macroeconomic mechanism that has so far been under-examined (Acemoglu & Autor, 2011). The results should therefore not be interpreted as a complete description of real economic dynamics, but as a theoretical frame of reference that should stimulate further empirical and model-theoretical work (Lucas, 1988; Baqaee & Farhi, 2019).

CONCLUSION AND OUTLOOK

This paper analyses the use of AI agents from a macroeconomic perspective and comes to a clear conclusion (Acemoglu & Restrepo, 2018). AI agents do not represent a gradual progress in productivity, but a structural break in the organisation of production and work (Aghion & Howitt, 1992; Acemoglu & Autor, 2011). Under competitive conditions, they lead to the complete substitution of entire task areas as soon as they can fulfil them more cost-effectively than human labour (Zeira, 1998; Acemoglu & Restrepo, 2018). In this context, automation is not an optional innovation path, but an endogenous equilibrium outcome (Romer, 1990; Aghion & Howitt, 1992). The developed task-based modelling framework makes this mechanism explicitly visible (Grossman & Rossi-Hansberg, 2008; Acemoglu & Autor, 2011). It shows that technological substitution cannot be explained by skill shifts or sectoral adjustment alone (Autor et al., 2003; Autor & Dorn, 2013). The decisive factor is the cost relationship between labour and AI agents (Grossman & Rossi-Hansberg, 2008; Zeira, 1998). Decreasing agent costs and increasing agent productivity systematically shift the automation threshold (Acemoglu & Restrepo, 2018). Competition ensures that this shift is effective across the board (Aghion & Howitt, 1992). Labour is not supplemented on automated tasks, but completely replaced (Zeira, 1998). The analysis makes it clear that productivity gains and employment development can permanently diverge in the context of AI agents (Acemoglu & Restrepo, 2020; Graetz & Michaels, 2018). Rising output does not necessarily go hand in hand with stable or growing labour demand (Basu & Fernald, 2000; Acemoglu & Restrepo, 2020). At the same time, the functional distribution of income is shifting from labour income to agent and capital income (Fuest et al., 2019; Bofinger, 2019). This shift is not a transitional phenomenon, but a structural consequence of the substitution of labour (Zeira, 1998). This results in potential risks for macroeconomic demand and macroeconomic stability (Summers, 2014). A key contribution of this paper is the explicit modelling of the competition-driven automation constraint (Aghion & Howitt, 1992; Grossman & Rossi-Hansberg, 2008). While much of the existing literature treats technological adoption as a gradual, partially discretionary process (Acemoglu & Autor, 2011), this work shows that AI agent automation is structurally enforced under competition (Zeira, 1998; Acemoglu & Restrepo, 2018). Adjustment costs and institutional inertia may delay the process, but do not change its long-term direction as long as the underlying cost relations persist (Acemoglu & Autor, 2011; Zeira, 1998). The results have far-reaching implications for theory, practice and economic policy. For macroeconomic theory, they suggest that central assumptions about the persistence of labour as a factor of production should be reconsidered (Solow, 1956; Lucas, 1988). For companies and markets, it becomes clear that automation by AI agents is less a strategic choice than a competitive necessity (Grossman & Rossi-Hansberg, 2008; Aghion & Howitt, 1992). For economic policy, the focus is shifting from preventing automation to shaping its macroeconomic consequences, particularly in terms of income distribution and demand protection (Summers, 2014; Truger, 2012). The outlook of this work opens up several research directions. A central task is to empirically test and calibrate the developed modelling framework (Baqaee & Farhi, 2019). In particular, the speed of automation, heterogeneity between sectors and the role of institutional barriers are empirically open questions (Acemoglu & Restrepo, 2020). In addition, an extension to explicit financial market and fiscal channels would be useful to analyse feedbacks between automation, asset prices and macroeconomic stability (Summers, 2014). Another important strand of research concerns the endogenisation of new tasks and activities (Acemoglu & Restrepo, 2018).

This paper deliberately does not treat task creation as an automatic equalisation mechanism (Zeira, 1998). Future work should investigate the conditions under which new tasks can emerge to a sufficient extent and whether they remain labour-intensive in the long term or can themselves be automated quickly (Acemoglu & Autor, 2011). Finally, the approach opens up a bridge to institutional and political economy (Snower, 1993). If automation is competition-driven and irreversible, questions of distribution, legitimisation and institutional design become more important (Bofinger, 2019; Truger, 2012). As a result, this work shows that AI agents mark a new phase of technological change (Acemoglu & Restrepo, 2018). They are not only changing how production takes place, but also the role of labour in modern economies (Zeira, 1998). This change is not a marginal phenomenon, but a fundamental macroeconomic process (Solow, 1956; Aghion & Howitt, 1992). A clear theoretical understanding of it is a necessary prerequisite for an objective and well-founded discussion of its economic and social consequences (Lucas, 1988; Summers, 2014).

Declaration by Authors

Acknowledgement: None

Source of Funding: None

Conflict of Interest: No conflicts of interest declared.

ABOUT AUTHOR

Enrico Moch studied business administration at the University of Wismar and subsequently completed a degree in business law at Saarland University. Several years later, he completed a doctorate in economics on the topic of the influence of artificial intelligence on the world of work. His main job is as an administrative officer at the DLR Project Management Agency in the Society, Innovation and Technology division, where he is assigned to the AI Applications in Business department. In this role, he deals with the monitoring, evaluation and structuring of research and innovation projects in the context of digital technologies. Enrico Moch also works as a university lecturer at various institutions in Germany, including the Baden-Württemberg Cooperative State University Ravensburg. He also holds the position of Assistant Professor of Economics at the IICUT. He has been Academic Director of the Herford-based GrandEdu Research School for several years. His research focuses on AI governance, technical data protection and the institutional governance of digital platforms. In terms of content, his work focuses on regulatory frameworks, governance structures and the economic impact of artificial intelligence and data-driven systems. Enrico Moch regularly publishes scientific articles, is involved in interdisciplinary book projects and is active in academic peer review. He is also active in knowledge transfer between academia and practice and hosts the podcast "GrandEdu Research School - On the trail of the economy!".

Dr. Peter Schlecht is Managing Director of GrandEdu GmbH, a leading continuing education institution that specialises in preparing professionals for degrees such as the Bachelor Professional, Master Professional and Master of Business Administration (MBA). In his role at GrandEdu, he is committed to developing innovative educational strategies that bridge the gap between vocational and academic education and contribute to the sustainable qualification of skilled workers in Germany. During his doctorate, Dr. Schlecht worked as a research assistant at the Department of Economics at the Niederrhein University of Applied Sciences. His research focused on the intersection of tax law, financial education and labour market policy, which today shapes his strategic work at GrandEdu.

REFERENCES

1. Acemoglu, D., & Autor, D. (2011). Chapter 12- Skills, tasks and technologies: Implications for employment and earnings. Handbook of Labor Economics, 4(Part B), 1043-1171. https://doi.org/10.1016/S0169-7218(11)02410-5
2. Acemoglu, D., & Restrepo, P. (2018). Artificial intelligence, automation and work (NBER Working Paper No. 24196). National Bureau of Economic Research. https://doi.org/10.3386/w24196
3. Acemoglu, D., & Restrepo, P. (2020). Robots and jobs: Evidence from U.S. labor markets. Journal of Political Economy, 128(6), 2188-2244. https://doi.org/10.1086/705716
4. Aghion, P., & Howitt, P. (1992). A model of growth through creative destruction.

Econometrica, 60(2), 323-351. https://doi.org/10.2307/2951599
5. Autor, D. H., & Dorn, D. (2013). The growth of low-skill service jobs and the polarization of the U.S. labor market. American Economic Review, 103(5), 1553-1597. https://doi.org/10.1257/aer.103.5.1553
6. Autor, D. H., Levy, F., & Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. Quarterly Journal of Economics, 118(4), 1279-1333. https://doi.org/10.1162/003355303322552801
7. Baqaee, D. R., & Farhi, E. (2019). The macroeconomic impact of microeconomic shocks: Beyond Hulten's theorem. Econometrica, 87(4), 1155-1203. https://doi.org/10.3982/ECTA15202
8. Basu, S., & Fernald, J. G. (2000). Why is productivity procyclical? Why do we care? (Working Paper No. WP-00-11). Federal Reserve Bank of Chicago.
9. Bofinger, P. (2019). Fundamentals of Economics: An Introduction to the Science of Markets (4th, updated edition). Pearson Studium.
10. Dixit, A. K., & Stiglitz, J. E. (1977). Monopolistic competition and optimum product diversity. American Economic Review, 67(3), 297–308.
11. Graetz, G., & Michaels, G. (2018). Robots at work. The Review of Economics and Statistics, 100(5), 753-768. https://doi.org/10.1162/rest_a_00754
12. Grossman, G. M., & Rossi-Hansberg, E. (2008). Trading tasks: A simple theory of offshoring. American Economic Review, 98(5), 1978-1997. https://doi.org/10.1257/aer.98.5.1978
13. Jones, C. I. (2011). Intermediate goods and weak links in the theory of economic development. American Economic Journal: Macroeconomics, 3(2), 1–28. https://doi.org/10.1257/mac.3.2.1
14. Lucas, R. E., Jr. (1988). On the mechanics of economic development. Journal of Monetary Economics, 22(1), 3-42. https://doi.org/10.1016/0304-3932(88)90168-7
15. Romer, P. M. (1990). Endogenous technological change. Journal of Political Economy, 98(5, Part 2), 71-102. https://doi.org/10.1086/261725
16. Schmidhuber, J. (2015). Deep learning in neural networks: An overview. Neural Networks, 61, 85–117. https://doi.org/10.1016/j.neunet.2014.09.003
17. Schularick, M., & Taylor, A. M. (2012). Credit booms gone bust: Monetary policy, leverage cycles, and financial crises, 1870–2008. American Economic Review, 102(2), 1029–1061. https://doi.org/10.1257/aer.102.2.1029
18. Snower, D. J. (1993). The future of the welfare state. The Economic Journal, 103(418), 700–717. https://doi.org/10.2307/2234539
19. Solow, R. M. (1956). A contribution to the theory of economic growth. The Quarterly Journal of Economics, 70(1), 65–94. https://doi.org/10.2307/1884513
20. Summers, L. H. (2014). U.S. economic prospects: Secular stagnation, hysteresis, and the zero lower bound. Business Economics, 49(2), 65–73. https://doi.org/10.1057/be.2014.13
21. Truger, A. (2012). Fiscal policy and employment in advanced and emerging economies: What can we learn from European experiences? (IMK Working Paper No. 92). Institut für Makroökonomie und Konjunkturforschung (IMK). https://www.imf.org/external/np/pp/eng/2012/061512.pdf
22. Zeira, J. (1998). Workers, machines, and economic growth. The Quarterly Journal of Economics, 113(4), 1091-1117. https://doi.org/10.1162/003355398555847

How to cite this article: Enrico Moch, Peter Schlecht. AI agents and the structural transformation of labour. *International Journal of Research and Review*. 2026; 13(1): 331-350. DOI: *https://doi.org/10.52403/ijrr.20260132*

******