Dr. Annemarie Kunn-Nelen
Associate professor
of Lifelong Learning

Prof. dr. Mark Levels
Program director ROA
University dean of Lifelong Learning

Deelsessie

1. Welke veranderingen beïnvloeden vandaag en morgen de arbeidsmarkt, en hoe?

2. Hoe krijgen we hier goed zicht op?

3. Is inzetten op een leven lang ontwikkelen wel écht de oplossing?

ROA

theguardian

April 18, 2018

Robots will take our jobs. We'd better plan now, before it's too late

ROA

Forbes

Sep 3, 2018

Will A Robot Take My Manufacturing Job? Yes, No, And Maybe

ROA

Een meerderheid van de werkgevers is bezig met het voorbereiden van automatisering...

99% bedrijven doet iets om voor te bereiden op de toekomst van werk (2019)

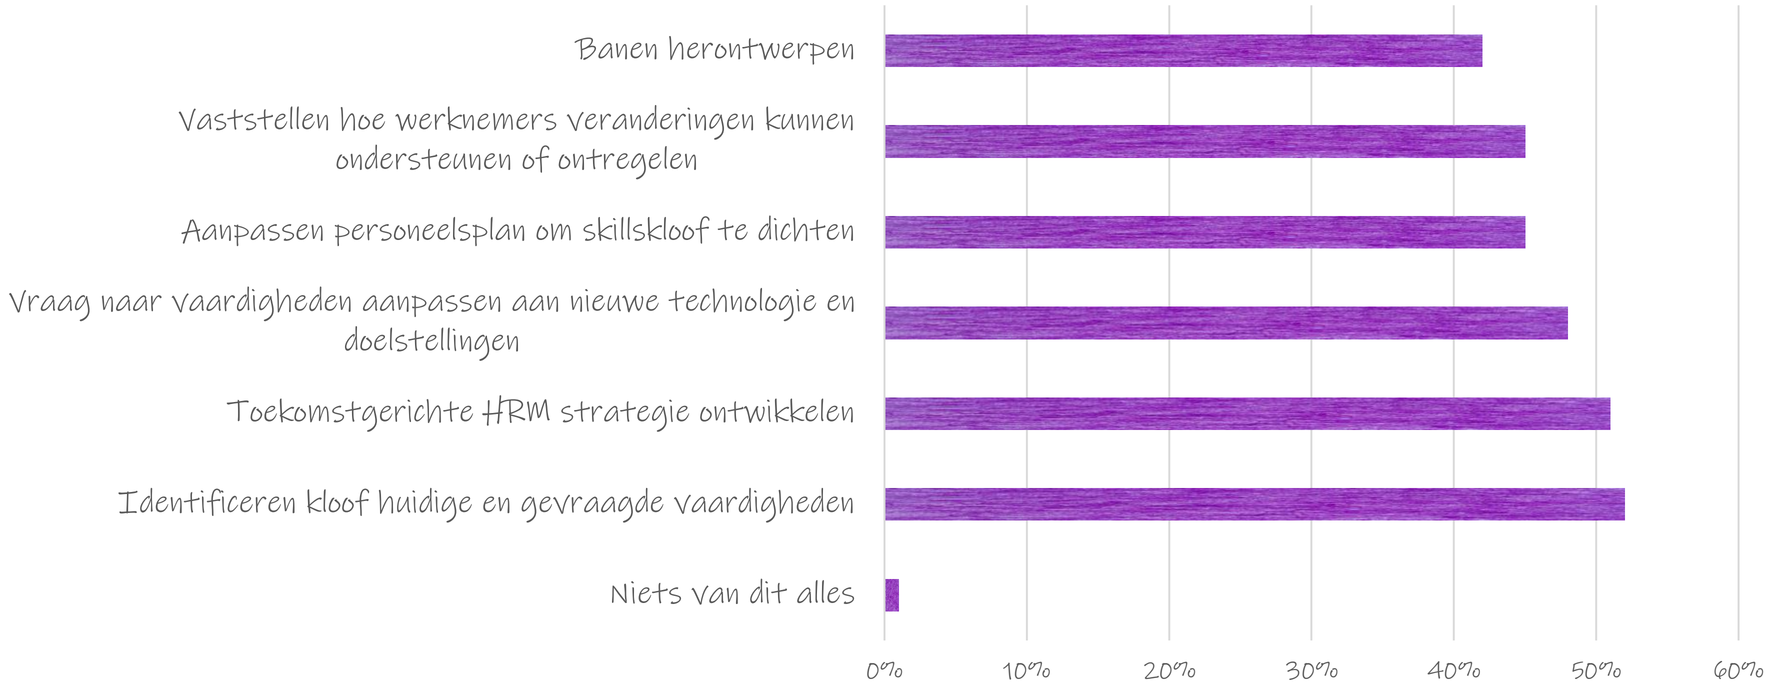

ROA

Bron: Mercer (2019) Global Telent Trends. Accessed on https://www.mercer.com/our-thinking/career/global-talent-hr-trends.html

UM

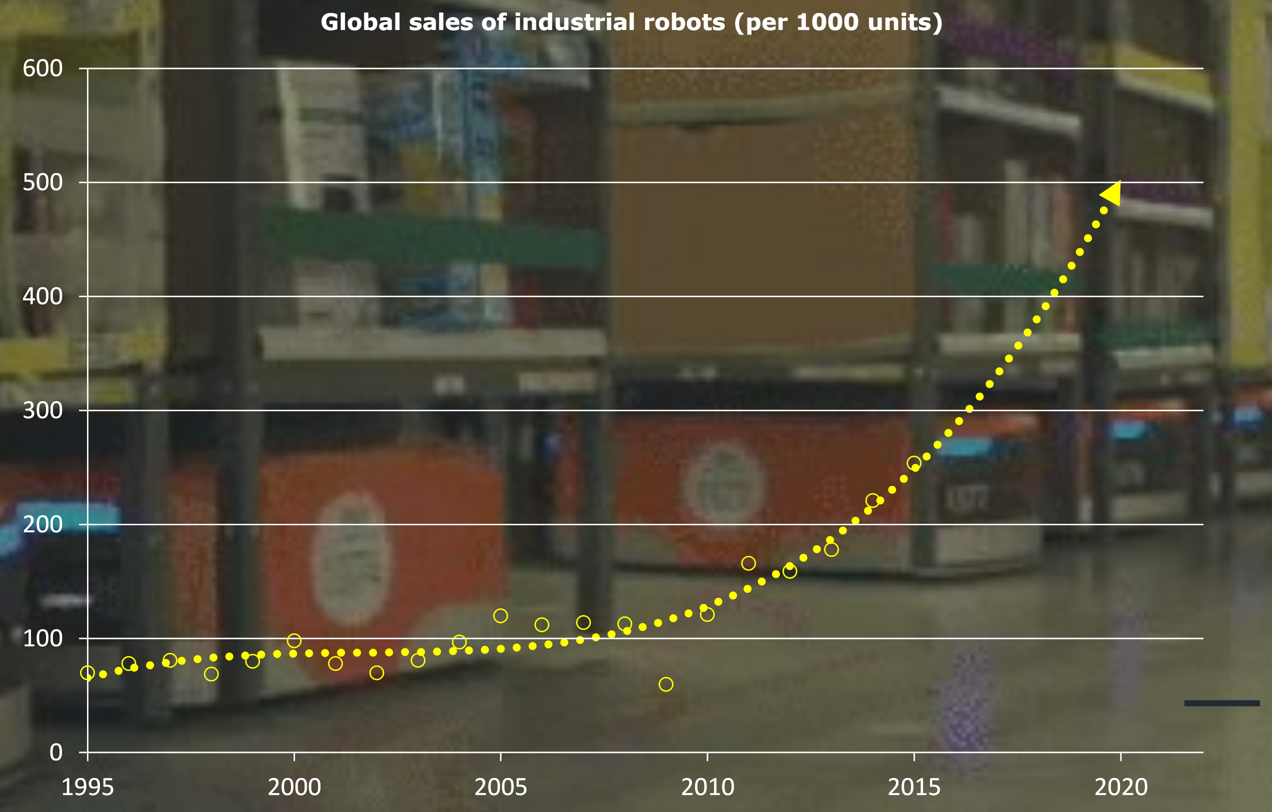

Source: International Federation of Robotics (2017), *World Robotics 2016 Industrial Robots, Frankfurt am Main: IFR*

Steeds minder technische flessenhalzen

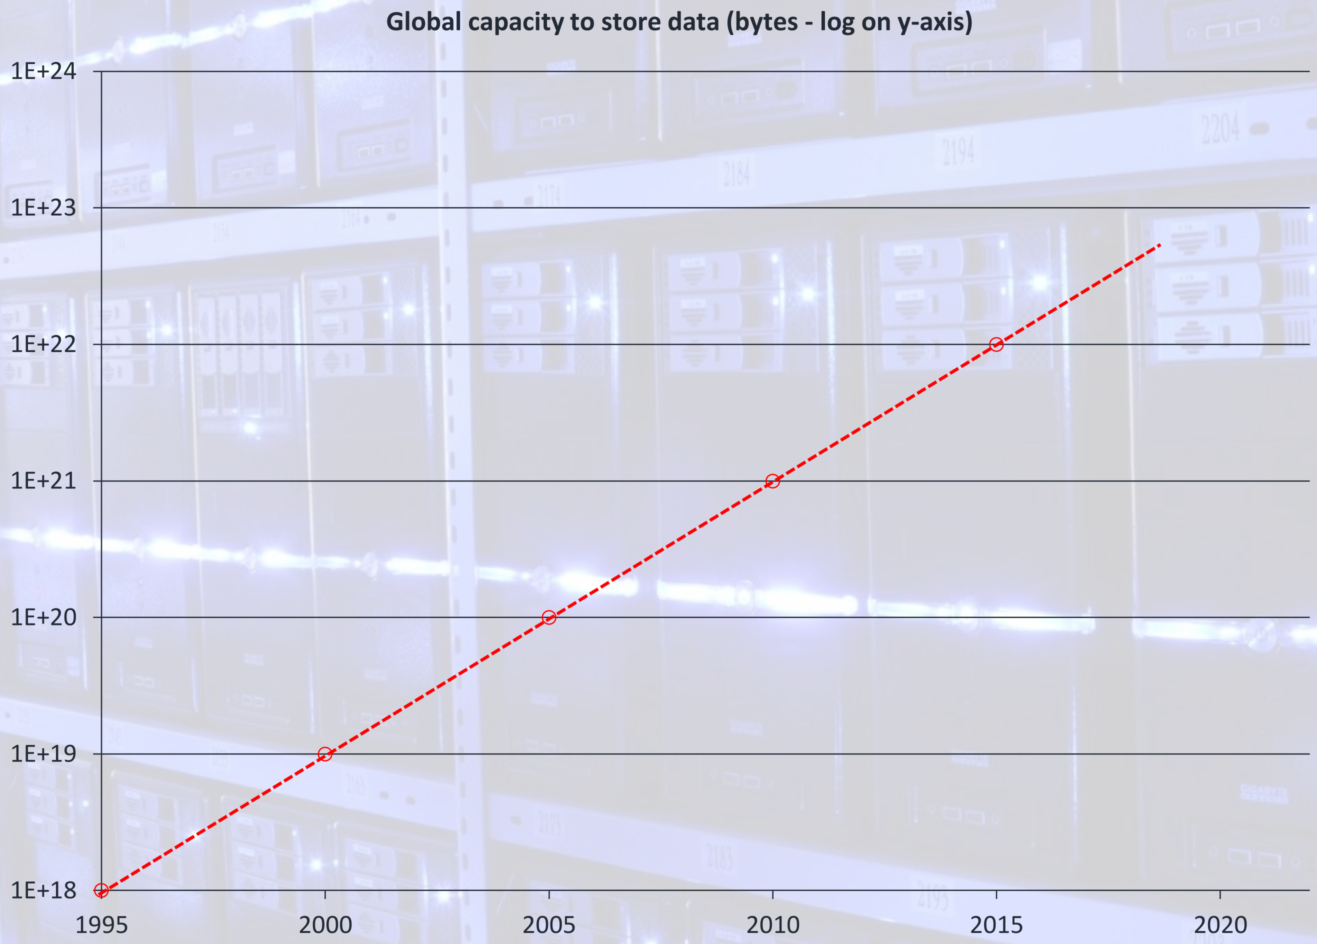

Source: Hilbert, M., and López, P. (2011). The World's Technological Capacity to Store, Communicate, and Compute Information. *Science*, 332(6025), 60 –65.

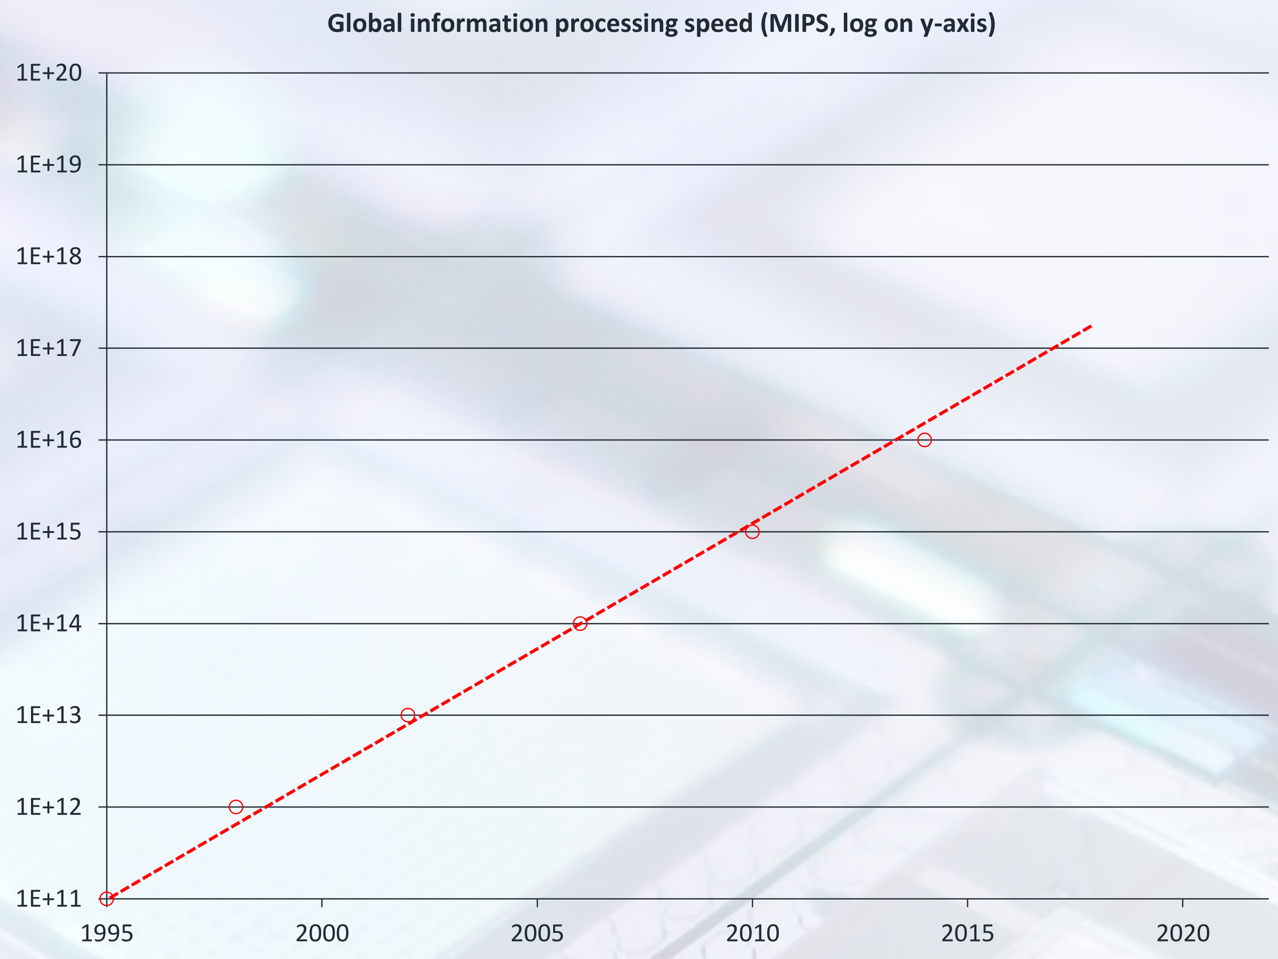

Source: Walldrop, M.M. (2016). More than Moore. *Nature* 530, 144–147.

Science RESEARCH ARTICLES

Cite as: N. Brown, T. Sandholm, *Science* 10.1126/science.aao1733 (2017).

Superhuman AI for heads-up no-limit poker: Libratus beats top professionals

Noam Brown and Tuomas Sandholm*

Computer Science Department, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA.
*Corresponding author. Email: sandholm@cs.cmu.edu

No-limit Texas hold'em is the most popular form of poker. Despite AI successes in perfect-information games, the private information and massive game tree have made no-limit poker difficult to tackle. We present Libratus, an AI that, in a 120,000-hand competition, defeated four top human specialist professionals in heads-up no-limit Texas hold'em, the leading benchmark and long-standing challenge problem in imperfect-information game solving. Our game-theoretic approach features application-independent techniques: an algorithm for computing a blueprint for the overall strategy, an algorithm that fleshes out the details of the strategy for subgames that are reached during play, and a self-improver algorithm that fixes potential weaknesses that opponents have identified in the blueprint strategy.

In recent years the field of artificial intelligence (AI) has advanced considerably. The measure of this progress has, in many cases, been marked by performance against humans in benchmark games. AI programs have defeated top humans in checkers (*1*), chess (*2*), and Go (*3*). In these perfect-information games both players know the exact state of the game at every point. In contrast, in imperfect-information games, some information about the state of the game is hidden from a player—for example, the opponent may hold hidden cards. Hidden information is ubiquitous in real-world strategic interactions, such as business strategy, negotiation, strategic pricing, finance, cybersecurity, and military applications, which makes research on general-purpose techniques for imperfect-information games particularly important.

Hidden information makes a game far more complex for a number of reasons. Rather than simply search for an optimal sequence of actions, an AI for imperfect-information games must determine how to balance actions appropriately, so that the opponent never finds out too much about the private information the AI has. For example, bluffing is a necessary feature in any competitive poker strategy, but bluffing all the time would be a bad strategy. In other words, the value of an action depends on the probability it is played.

Another key challenge is that different parts of the game cannot be considered in isolation; the optimal strategy for a given situation may depend on the strategy that would be played in situations that have not occurred (*4*). As a conse-

world. The heads-up (that is, two-player) variant prevents opponent collusion and kingmaker scenarios where a bad player causes a mediocre player to shine, and therefore allows a clear winner to be determined. Due to its large size and strategic complexity, heads-up no-limit Texas hold'em (HUNL) has been the primary benchmark and challenge problem for imperfect-information game solving for several years. No prior AI has defeated top human players in this game.

In this paper we introduce Libratus, (*12*) an AI that takes a distinct approach to addressing imperfect-information games. In a 20-day, 120,000-hand competition featuring a \$200,000 prize pool, it defeated top human professionals in HUNL. The techniques in Libratus do not use expert domain knowledge or human data and are not specific to poker; thus they apply to a host of imperfect-information games.

Game-solving approach in Libratus

Libratus features three main modules:

(i) The first module computes an abstraction of the game, which is smaller and easier to solve, and then computes game-theoretic strategies for the abstraction. The solution to this abstraction provides a detailed strategy for the early rounds of the game, but only an approximation for how to play in the more numerous later parts of the game. We refer to the solution of the abstraction as the blueprint strategy.

(ii) When a later part of the game is reached during play,

LETTER

nature

doi:10.1038/nature14236

Human-level control through deep reinforcement learning

Volodymyr Mnih[1*], Koray Kavukcuoglu[1*], David Silver[1*], Andrei A. Rusu[1], Joel Veness[1], Marc G. Bellemare[1], Alex Graves[1], Martin Riedmiller[1], Andreas K. Fidjeland[1], Georg Ostrovski[1], Stig Petersen[1], Charles Beattie[1], Amir Sadik[1], Ioannis Antonoglou[1], Helen King[1], Dharshan Kumaran[1], Daan Wierstra[1], Shane Legg[1] & Demis Hassabis[1]

The theory of reinforcement learning provides a normative account[1], deeply rooted in psychological[2] and neuroscientific[3] perspectives on animal behaviour, of how agents may optimize their control of an environment. To use reinforcement learning successfully in situations approaching real-world complexity, however, agents are confronted with a difficult task: they must derive efficient representations of the environment from high-dimensional sensory inputs, and use these to generalize past experience to new situations. Remarkably, humans and other animals seem to solve this problem through a harmonious combination of reinforcement learning and hierarchical sensory processing systems[4,5], the former evidenced by a wealth of neural data revealing notable parallels between the phasic signals emitted by dopaminergic neurons and temporal difference reinforcement learning algorithms[3]. While reinforcement learning agents have achieved some successes in a variety of domains[6–8], their applicability has previously been limited to domains in which useful features can be handcrafted, or to domains with fully observed, low-dimensional state spaces. Here we use recent advances in training deep neural networks[9–11] to develop a novel artificial agent, termed a deep Q-network, that can learn successful policies directly from high-dimensional sensory inputs using end-to-end reinforcement learning. We tested this agent on the challenging domain of classic Atari 2600 games[12]. We demonstrate that the deep Q-network agent, receiving only the pixels and the game score as inputs, was able to surpass the performance of all previous algorithms and achieve a level comparable to that of a professional human games tester across a set of 49 games, using the same algorithm, network architecture and hyperparameters. This work bridges the divide between high-dimensional sensory inputs and actions, resulting in the first artificial agent that is capable of learning to excel at a diverse array of challenging tasks.

We set out to create a single algorithm that would be able to develop a wide range of competencies on a varied range of challenging tasks—a central goal of general artificial intelligence[13] that has eluded previous efforts[8,14,15]. To achieve this, we developed a novel agent, a deep Q-network (DQN), which is able to combine reinforcement learning with a class of artificial neural network[16] known as deep neural networks. Notably, recent advances in deep neural networks[9–11], in which several layers of nodes are used to build up progressively more abstract representations of the data, have made it possible for artificial neural networks to learn concepts such as object categories directly from raw sensory data. We use one particularly successful architecture, the deep convolutional network[17], which uses hierarchical layers of tiled convolutional filters to mimic the effects of receptive fields—inspired by Hubel and Wiesel's seminal work on feedforward processing in early visual cortex[18]—thereby exploiting the local spatial correlations present in images, and building in robustness to natural transformations such as changes of viewpoint or scale.

We consider tasks in which the agent interacts with an environment through a sequence of observations, actions and rewards. The goal of the agent is to select actions in a fashion that maximizes cumulative future reward. More formally, we use a deep convolutional neural network to approximate the optimal action-value function

$$Q^*(s,a) = \max_{\pi} \mathbb{E}\left[r_t + \gamma r_{t+1} + \gamma^2 r_{t+2} + \ldots \middle| s_t = s,\ a_t = a,\ \pi\right],$$

which is the maximum sum of rewards r_t discounted by γ at each time-step t, achievable by a behaviour policy $\pi = P(a|s)$, after making an observation (s) and taking an action (a) (see Methods)[19].

Reinforcement learning is known to be unstable or even to diverge when a nonlinear function approximator such as a neural network is used to represent the action-value (also known as Q) function[20]. This instability has several causes: the correlations present in the sequence of observations, the fact that small updates to Q may significantly change the policy and therefore change the data distribution, and the correlations between the action-values (Q) and the target values $r + \gamma \max_{a'} Q(s', a')$. We address these instabilities with a novel variant of Q-learning, which uses two key ideas. First, we used a biologically inspired mechanism termed experience replay[21–23] that randomizes over the data, thereby removing correlations in the observation sequence and smoothing over changes in the data distribution (see below for details). Second, we used an iterative update that adjusts the action-values (Q) towards target values that are only periodically updated, thereby reducing correlations with the target.

While other stable methods exist for training neural networks in the reinforcement learning setting, such as neural fitted Q-iteration[24], these methods involve the repeated training of networks *de novo* on hundreds of iterations. Consequently, these methods, unlike our algorithm, are too inefficient to be used successfully with large neural networks. We parameterize an approximate value function $Q(s,a;\theta_i)$ using the deep convolutional neural network shown in Fig. 1, in which θ_i are the parameters (that is, weights) of the Q-network at iteration i. To perform experience replay we store the agent's experiences $e_t = (s_t,a_t,r_t,s_{t+1})$ at each time-step t in a data set $D_t = \{e_1,\ldots,e_t\}$. During learning, we apply Q-learning updates, on samples (or minibatches) of experience $(s,a,r,s') \sim U(D)$, drawn uniformly at random from the pool of stored samples. The Q-learning update at iteration i uses the following loss function:

$$L_i(\theta_i) = \mathbb{E}_{(s,a,r,s') \sim \mathrm{U}(D)}\left[\left(r + \gamma \max_{a'} Q(s',a';\theta_i^-) - Q(s,a;\theta_i)\right)^2\right]$$

in which γ is the discount factor determining the agent's horizon, θ_i are the parameters of the Q-network at iteration i and θ_i^- are the network parameters used to compute the target at iteration i. The target network parameters θ_i^- are only updated with the Q-network parameters (θ_i) every C steps and are held fixed between individual updates (see Methods).

To evaluate our DQN agent, we took advantage of the Atari 2600 platform, which offers a diverse array of tasks ($n = 49$) designed to be

[1]Google DeepMind, 5 New Street Square, London EC4A 3TW, UK.
*These authors contributed equally to this work.

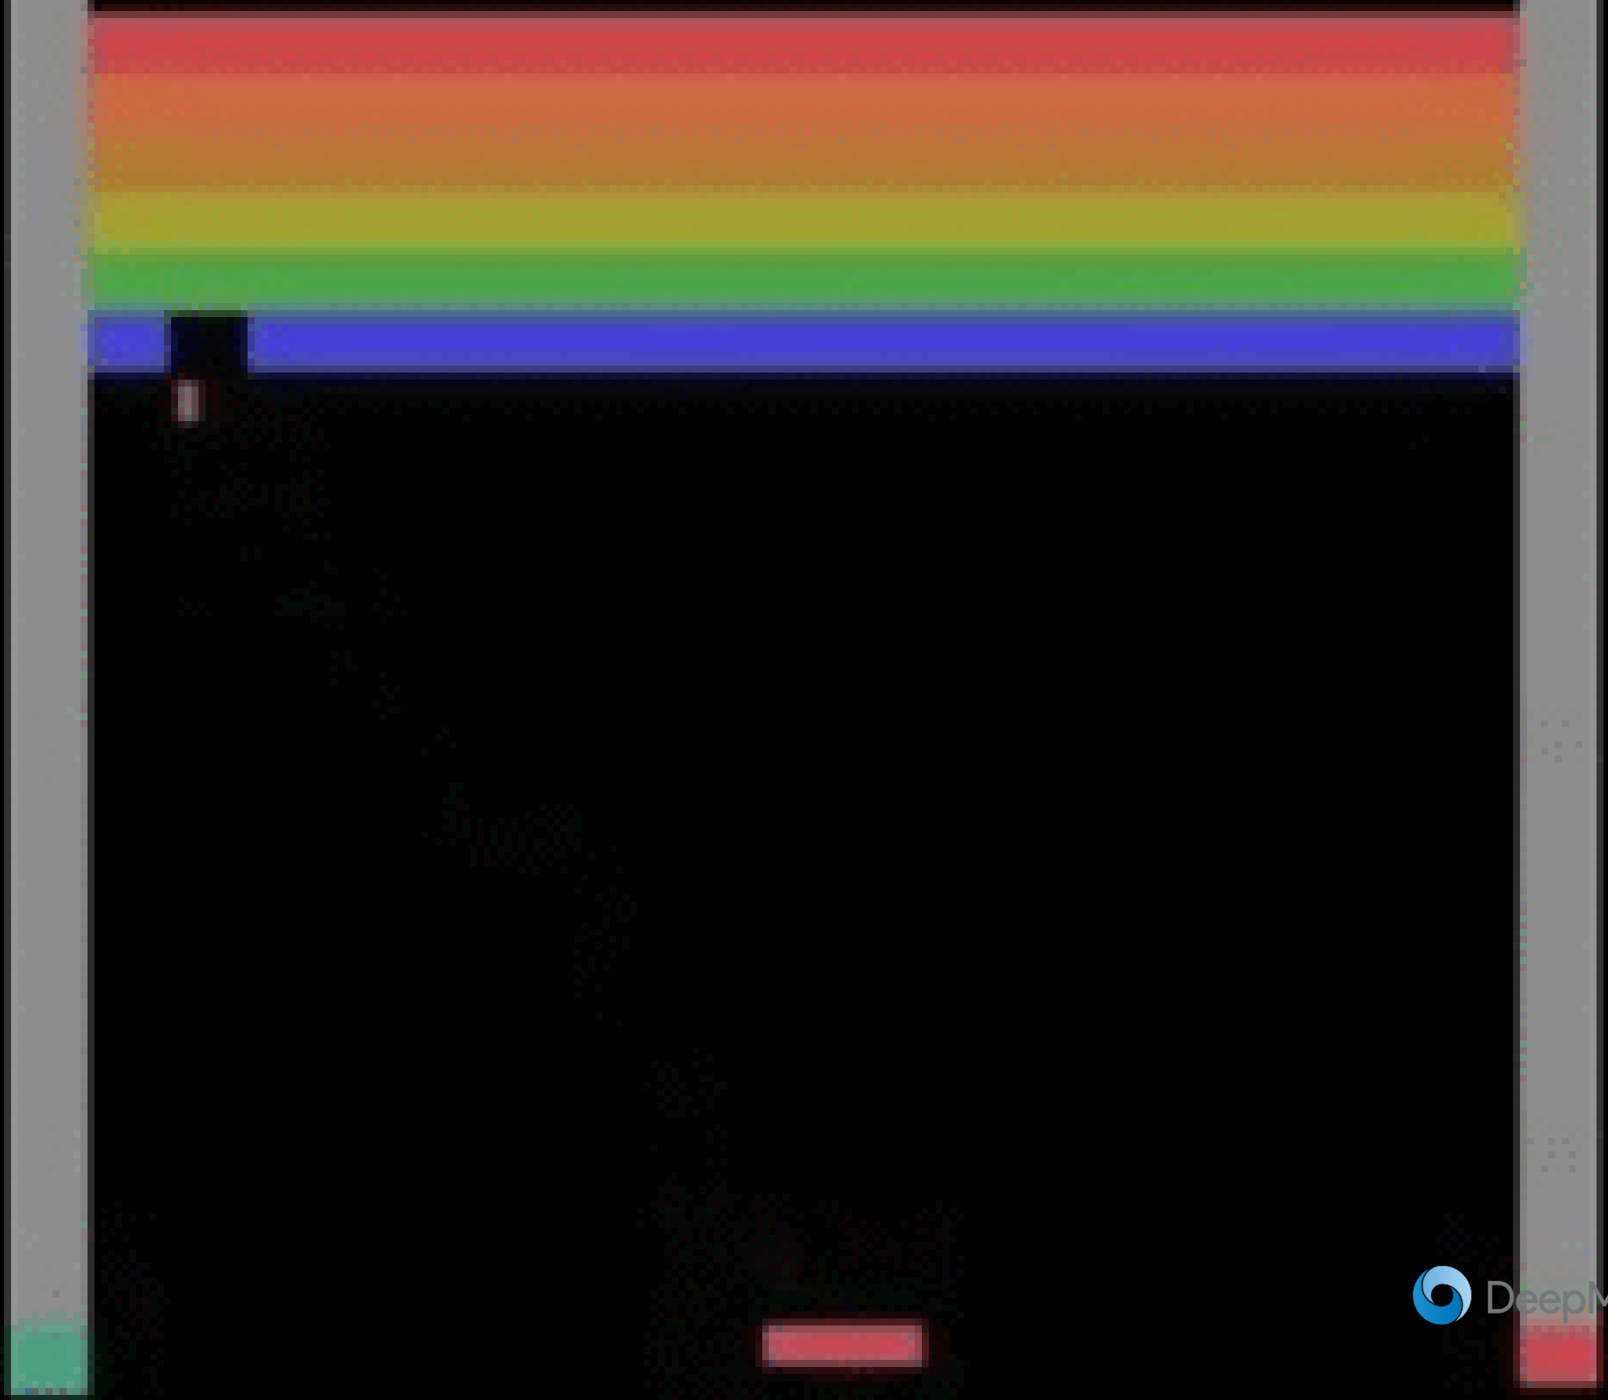

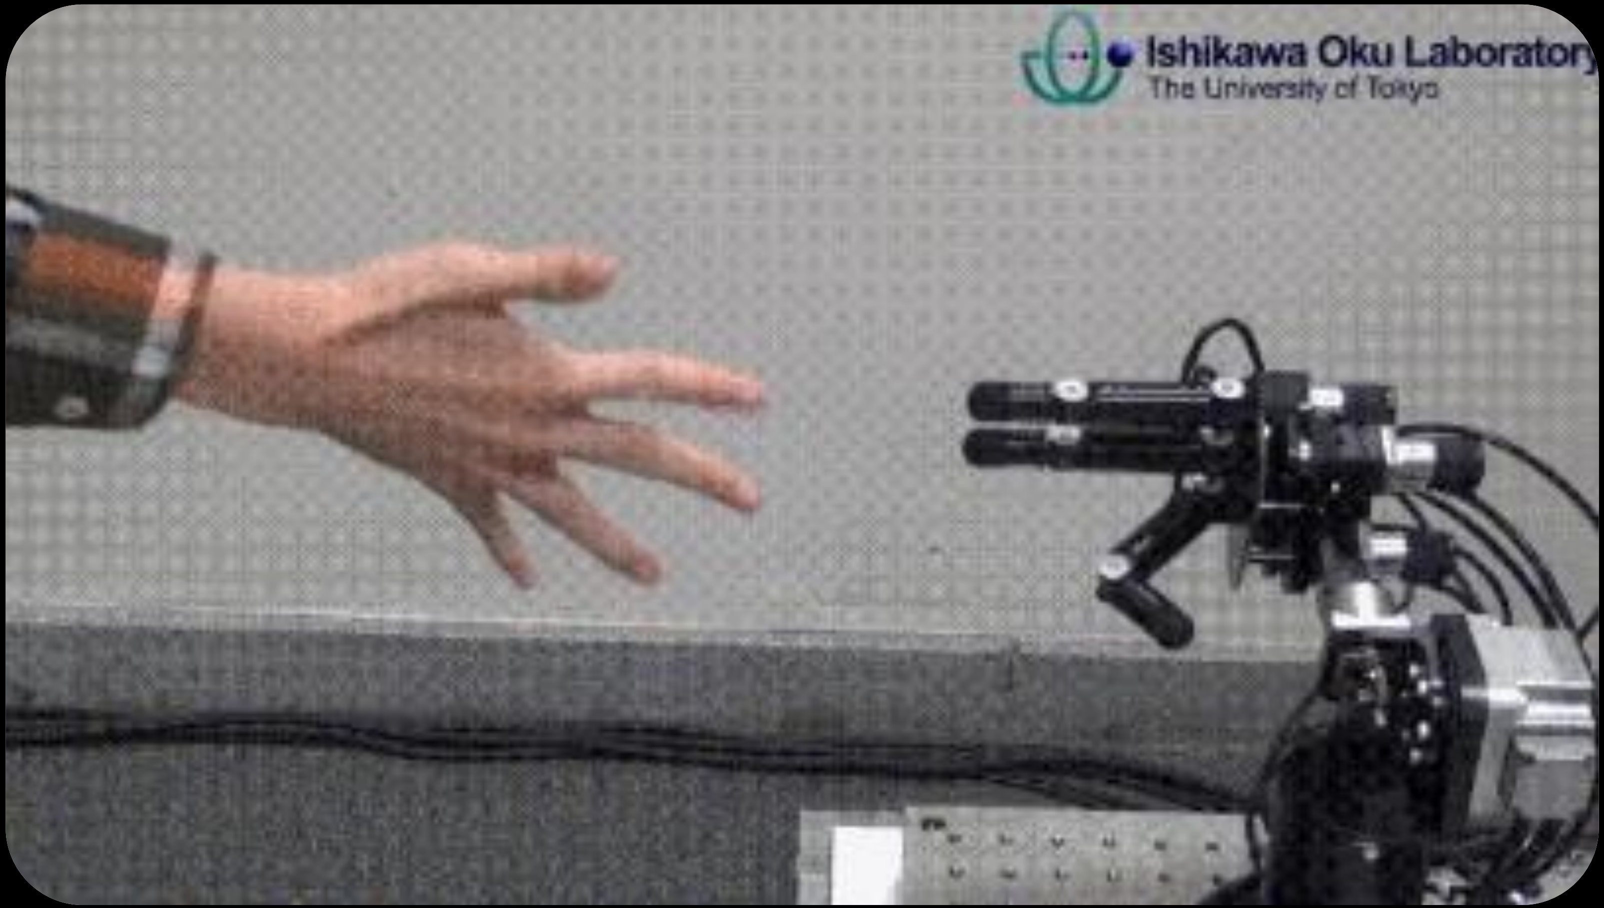
Ishikawa Oku Laboratory
The University of Tokyo

The New York Times
April 18, 2018

Robot Conquers One of the Hardest Human Tasks: Assembling Ikea Furniture

Wat betekent dit voor arbeidsmarkt en samenleving?

THE NEW YORK TIMES, SUNDAY, FEBRUARY 26, 1928. XX 3

MARCH OF THE MACHINE MAKES IDLE HANDS

Prevalence of Unemployment With Greatly Increased Industrial Output Points to the Influence of Labor-Saving Devices as an Underlying Cause

By EVANS CLARK.

A FEW days ago the General Motors Corporation reported the largest peace-time earnings ever made by a single concern in the history of America. Three days later Governor Smith made public a report from the New York Industrial Commissioner which called public attention to serious unemployment throughout the State: not since the depression of 1921, it was disclosed, have conditions been as bad.

The people of the United States—in the shadow of a Presidential election—are presented with a social

have gone far to make construction a machine industry instead of a collection of hand trades. One gasoline crane takes the place of ten or twelve laborers. The hod-carrier has disappeared before the invasion of the material hoist. In concrete construction building materials are mixed, like dough, in a machine and literally poured into place without the touch of a human hand. The Ohio figures record these results: with 15 per cent. fewer men employed, contractors put up 11 per cent. more square feet of finished buildings last year than in 1923.

Coal Mined by Machines.

ROA

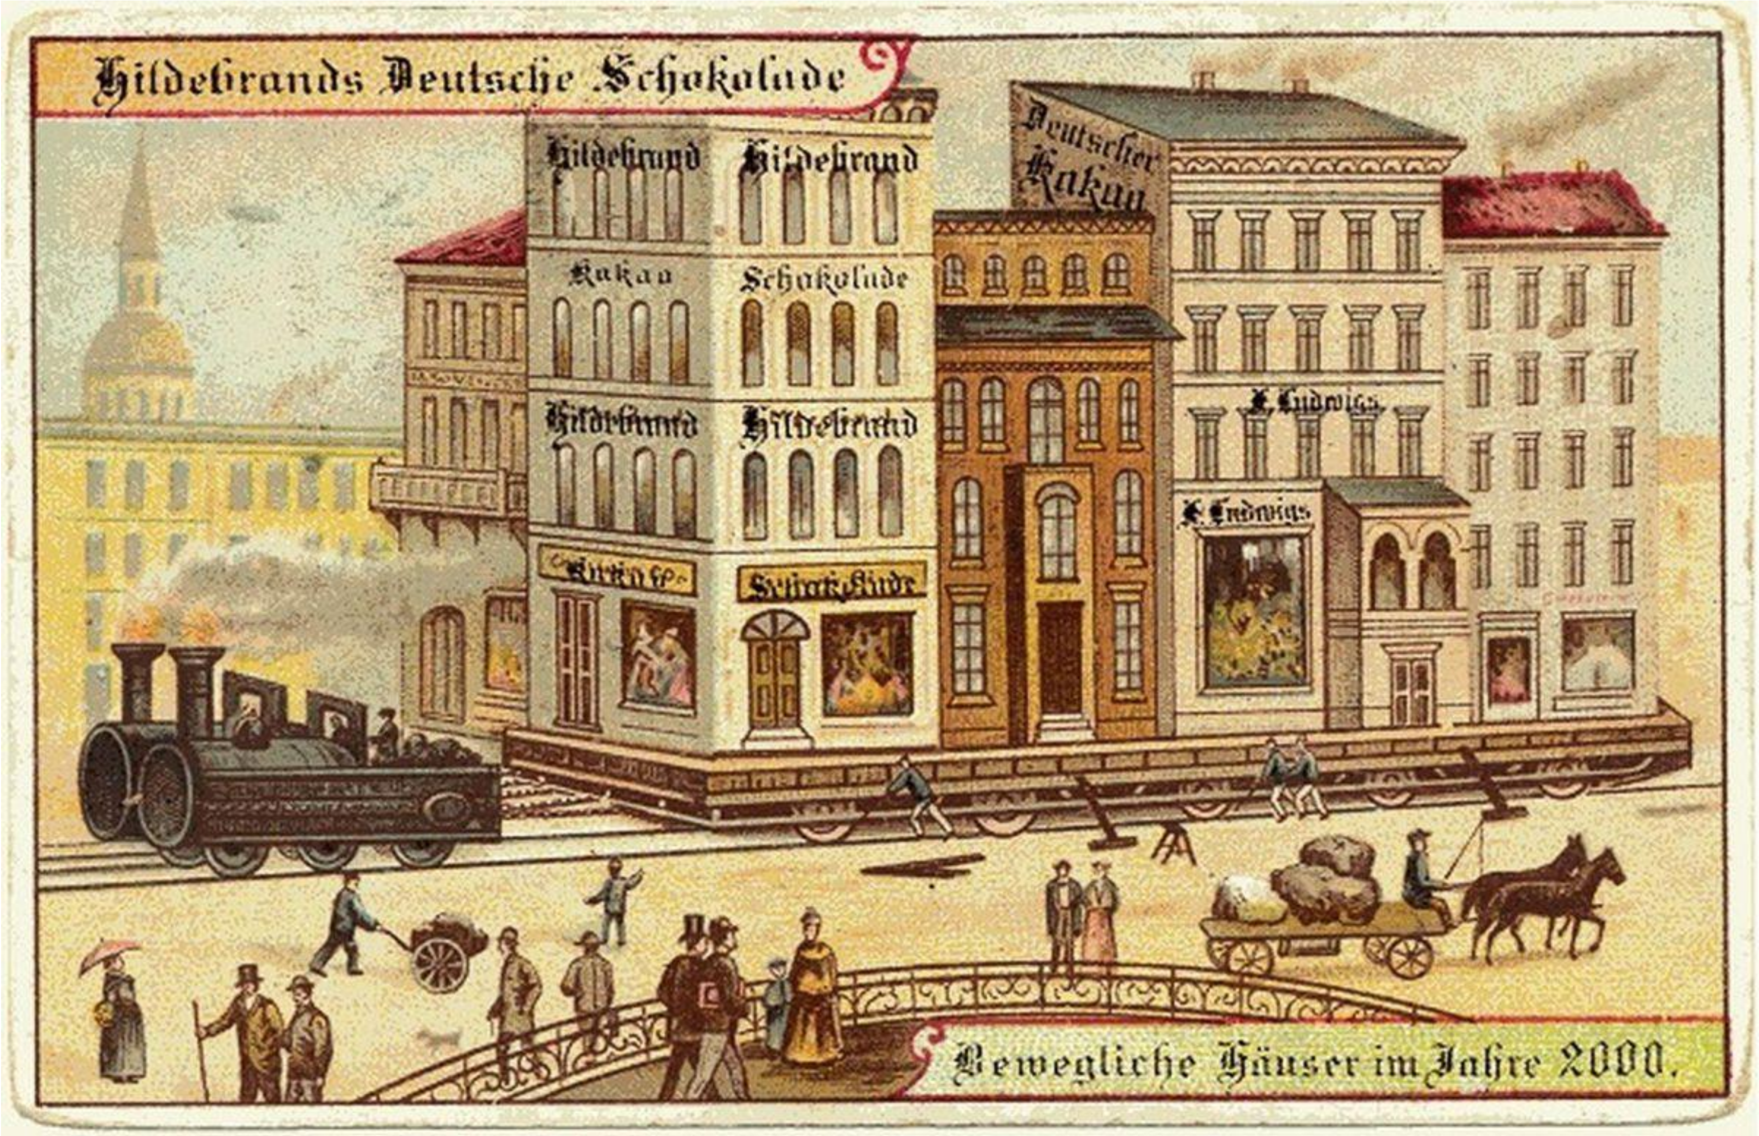
Hildebrands Deutsche Schokolade
Hildebrand
Hildebrand
Kakao
Schokolade
Deutscher Kakao
Bewegliche Häuser im Jahre 2000.

ROA

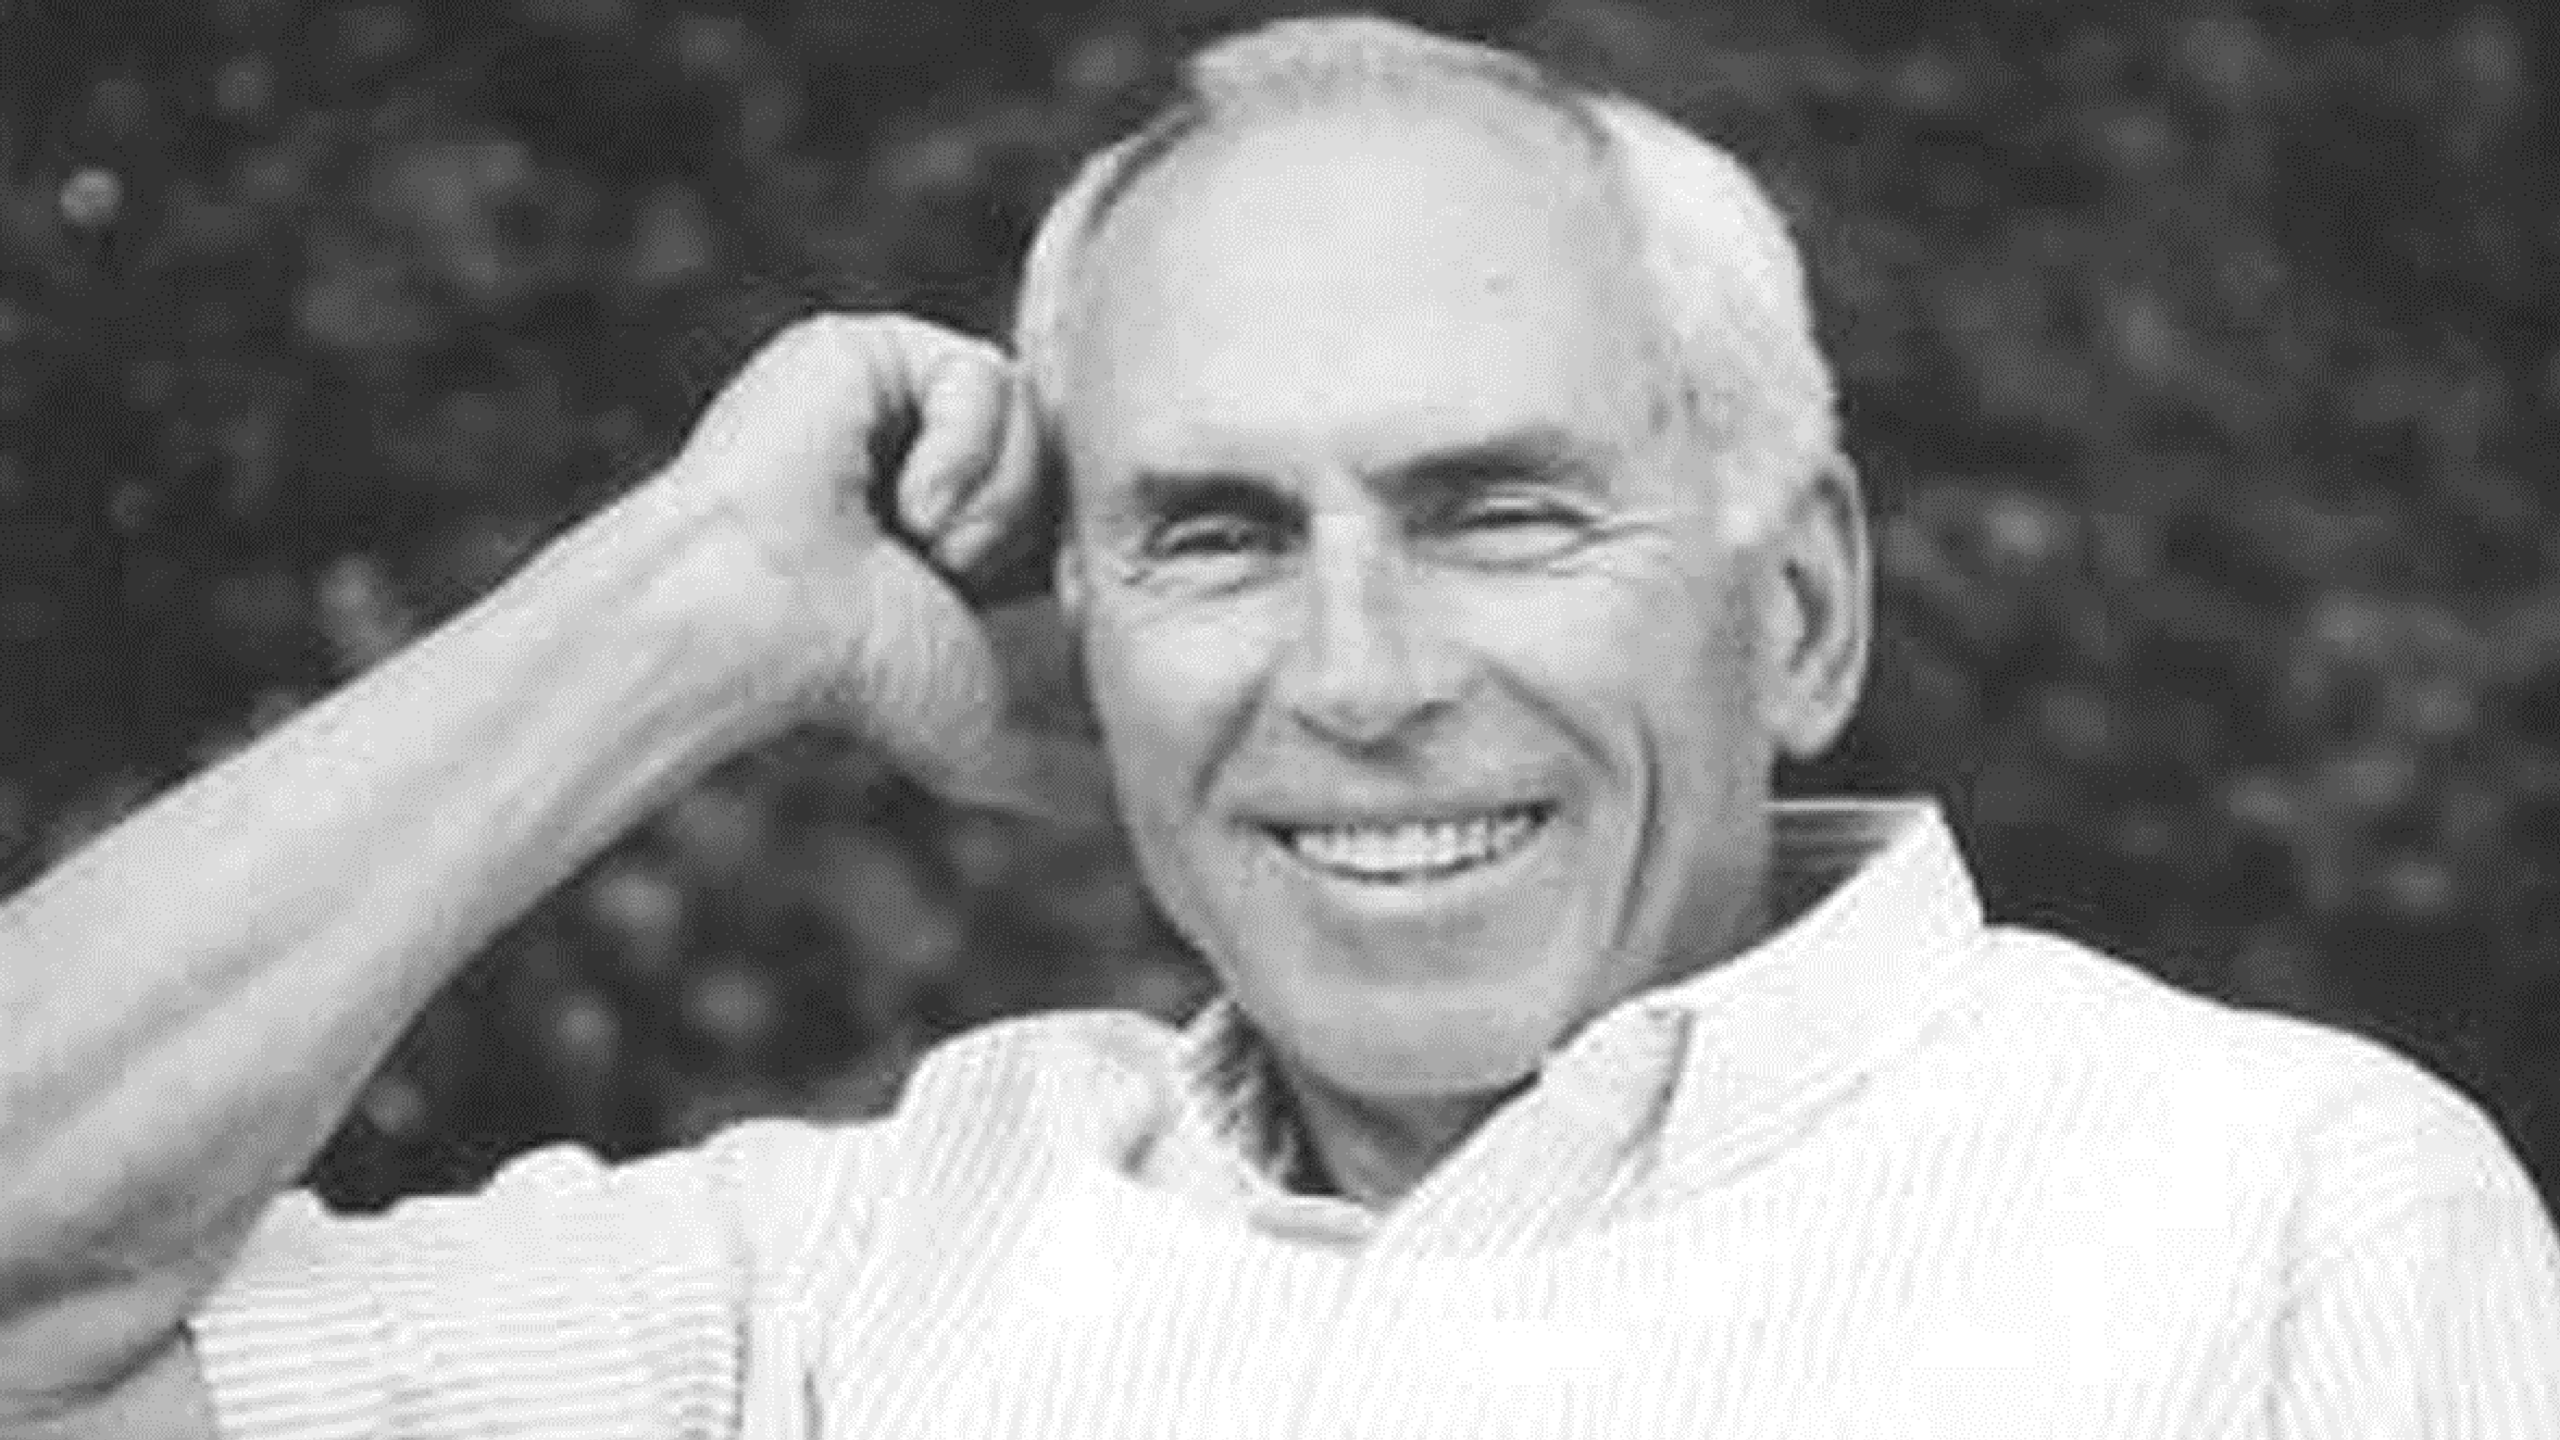

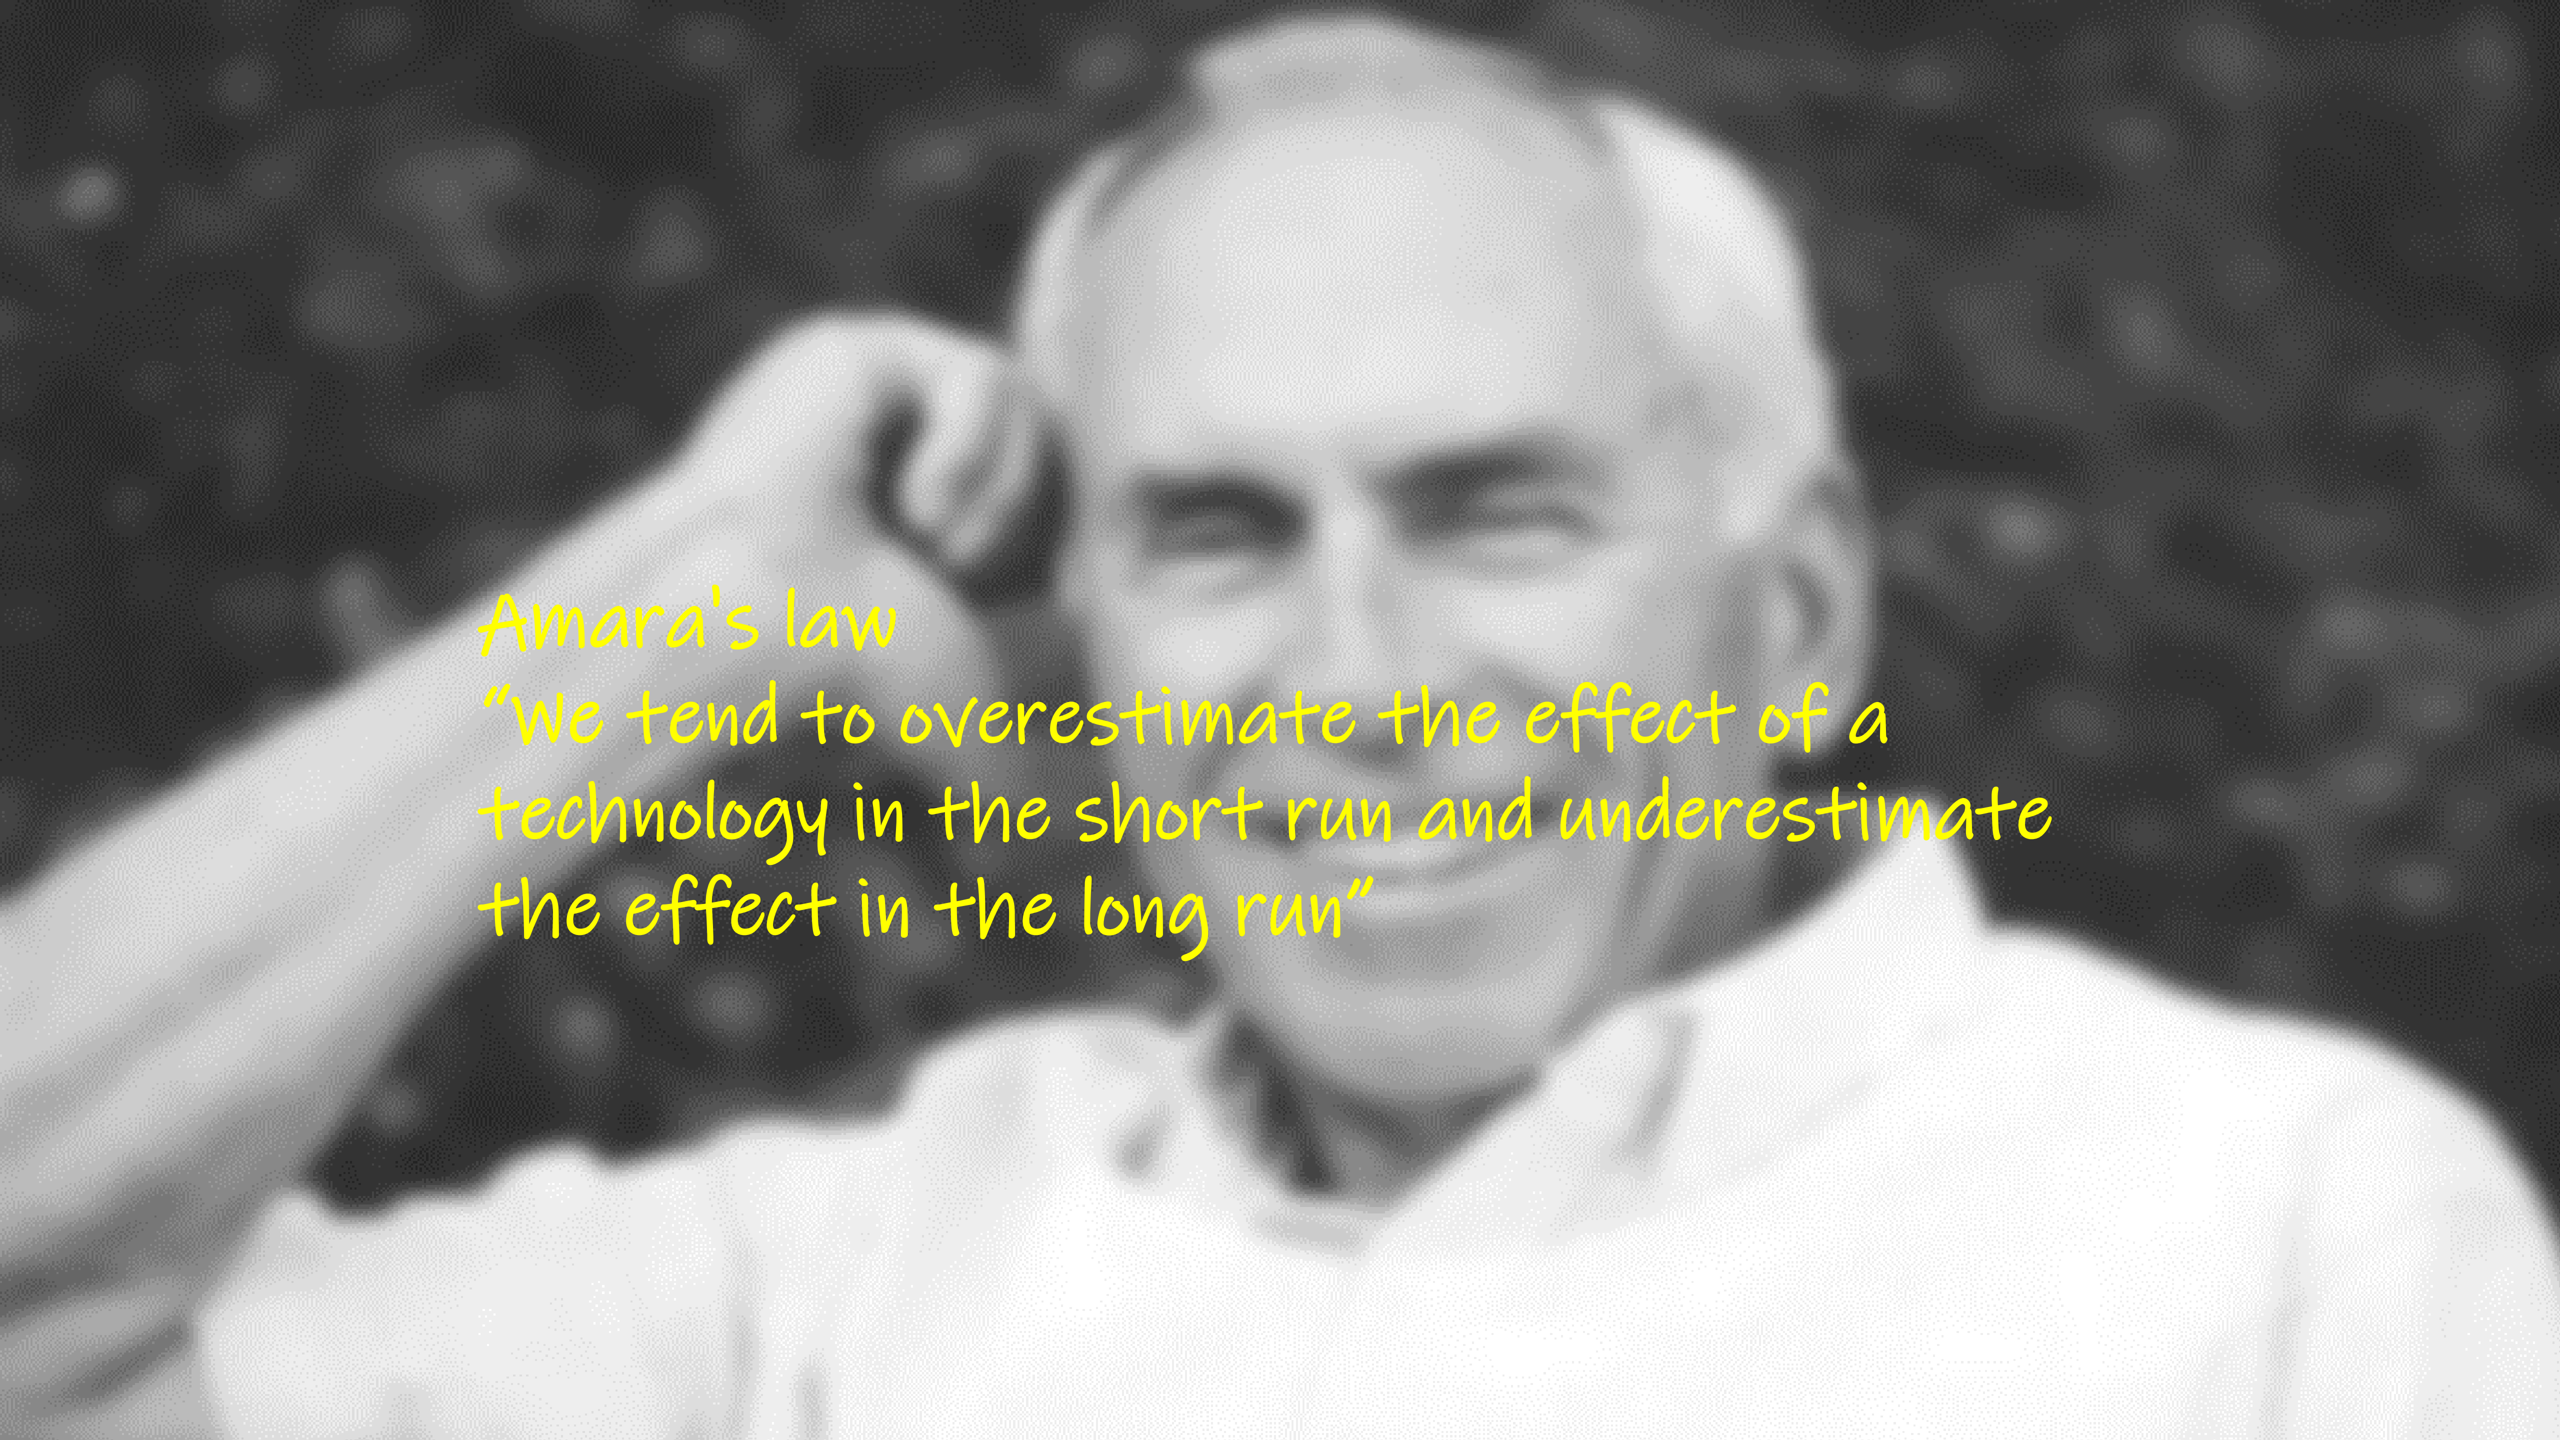
Amara's law
"We tend to overestimate the effect of a technology in the short run and underestimate the effect in the long run"


iRobot

Wetenschappelijk onderzoek!

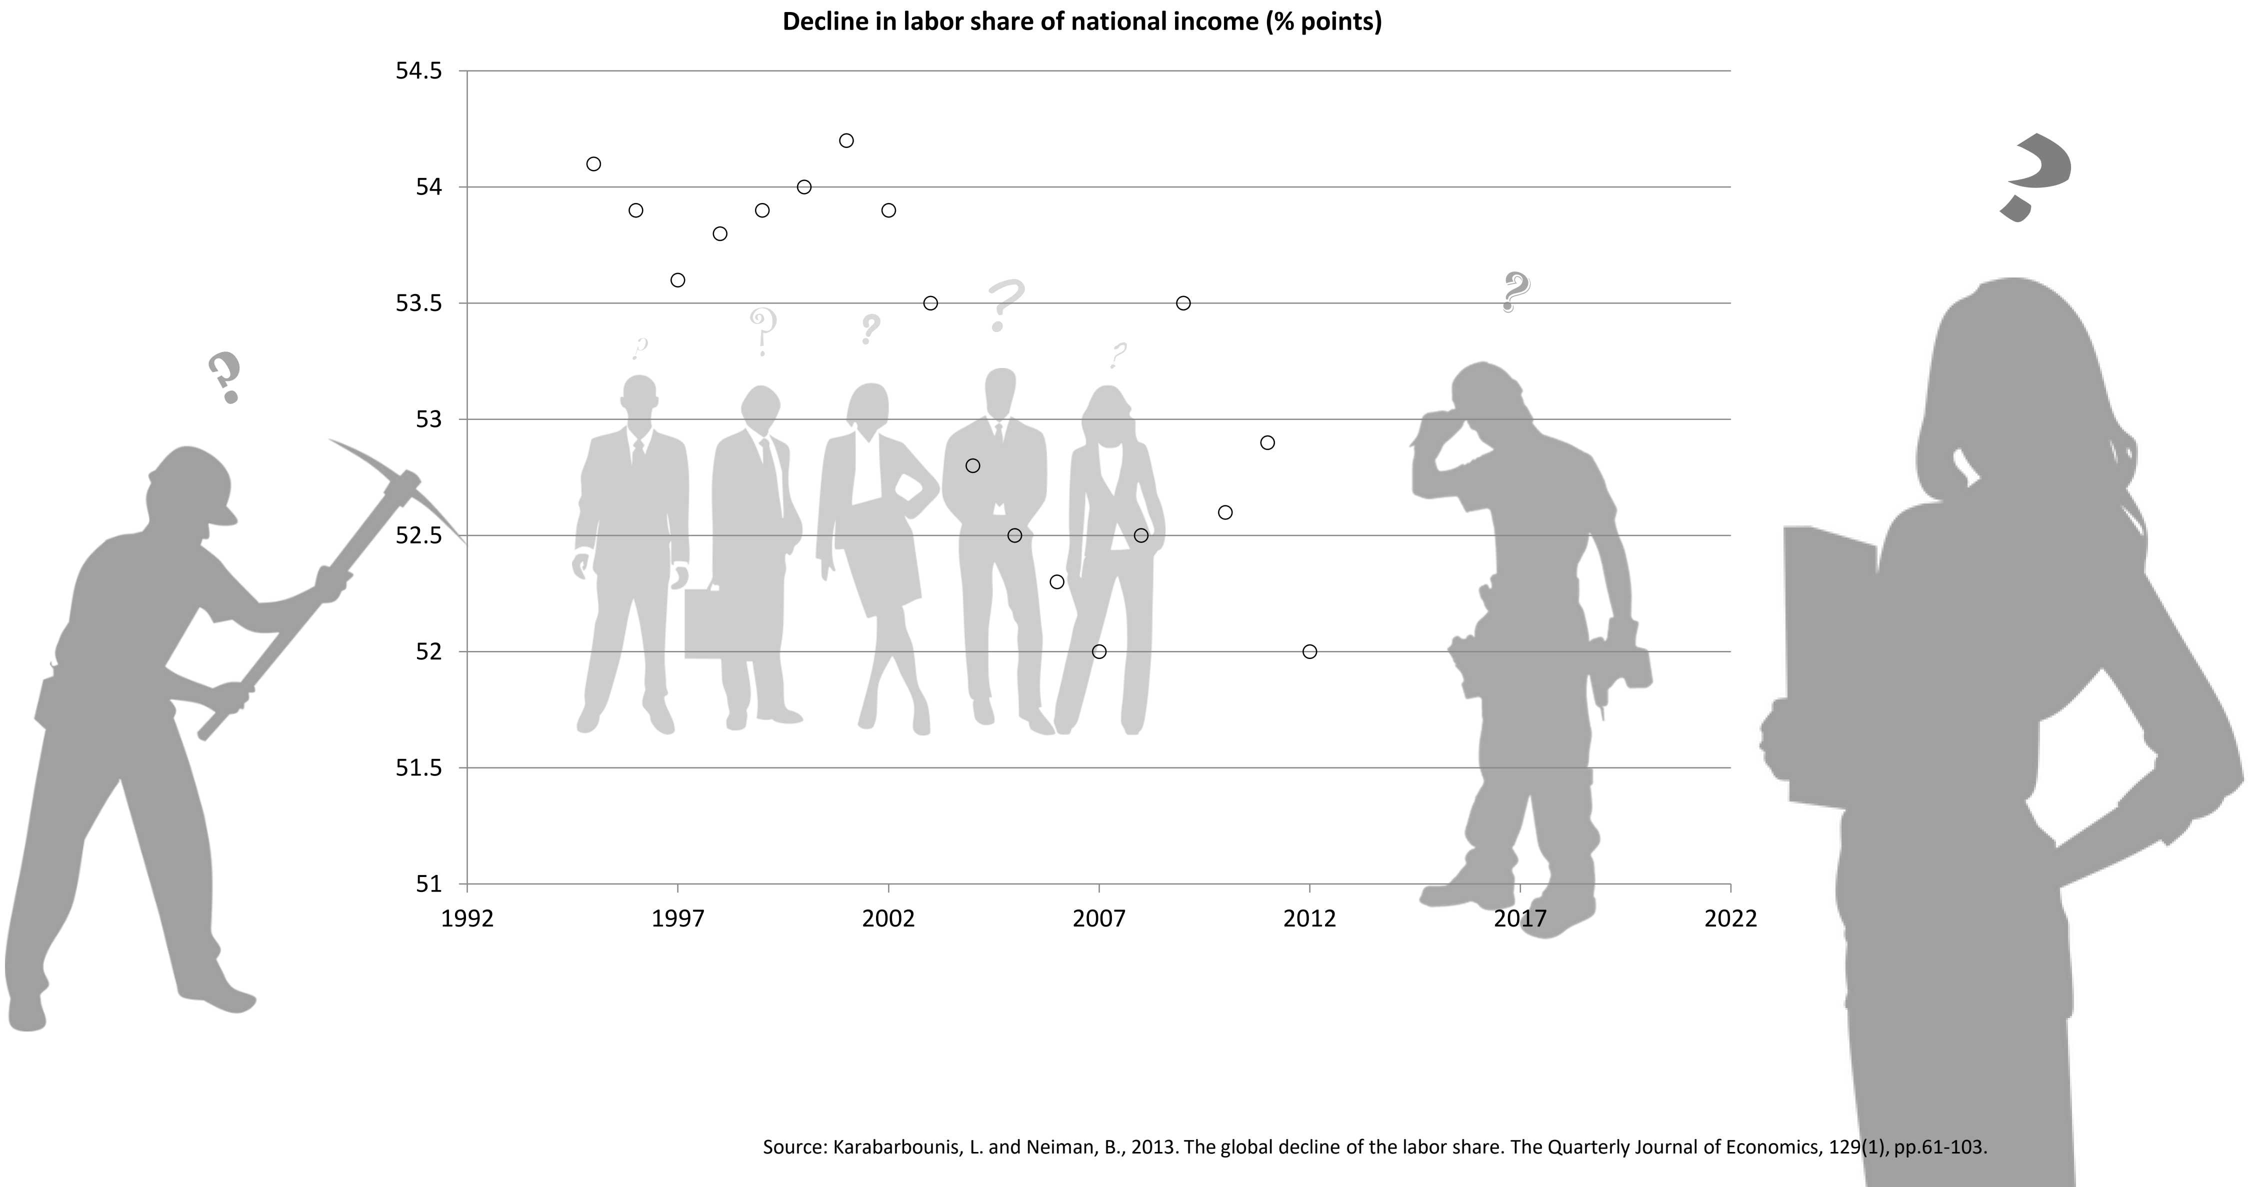

Source: Karabarbounis, L. and Neiman, B., 2013. The global decline of the labor share. The Quarterly Journal of Economics, 129(1), pp.61-103.

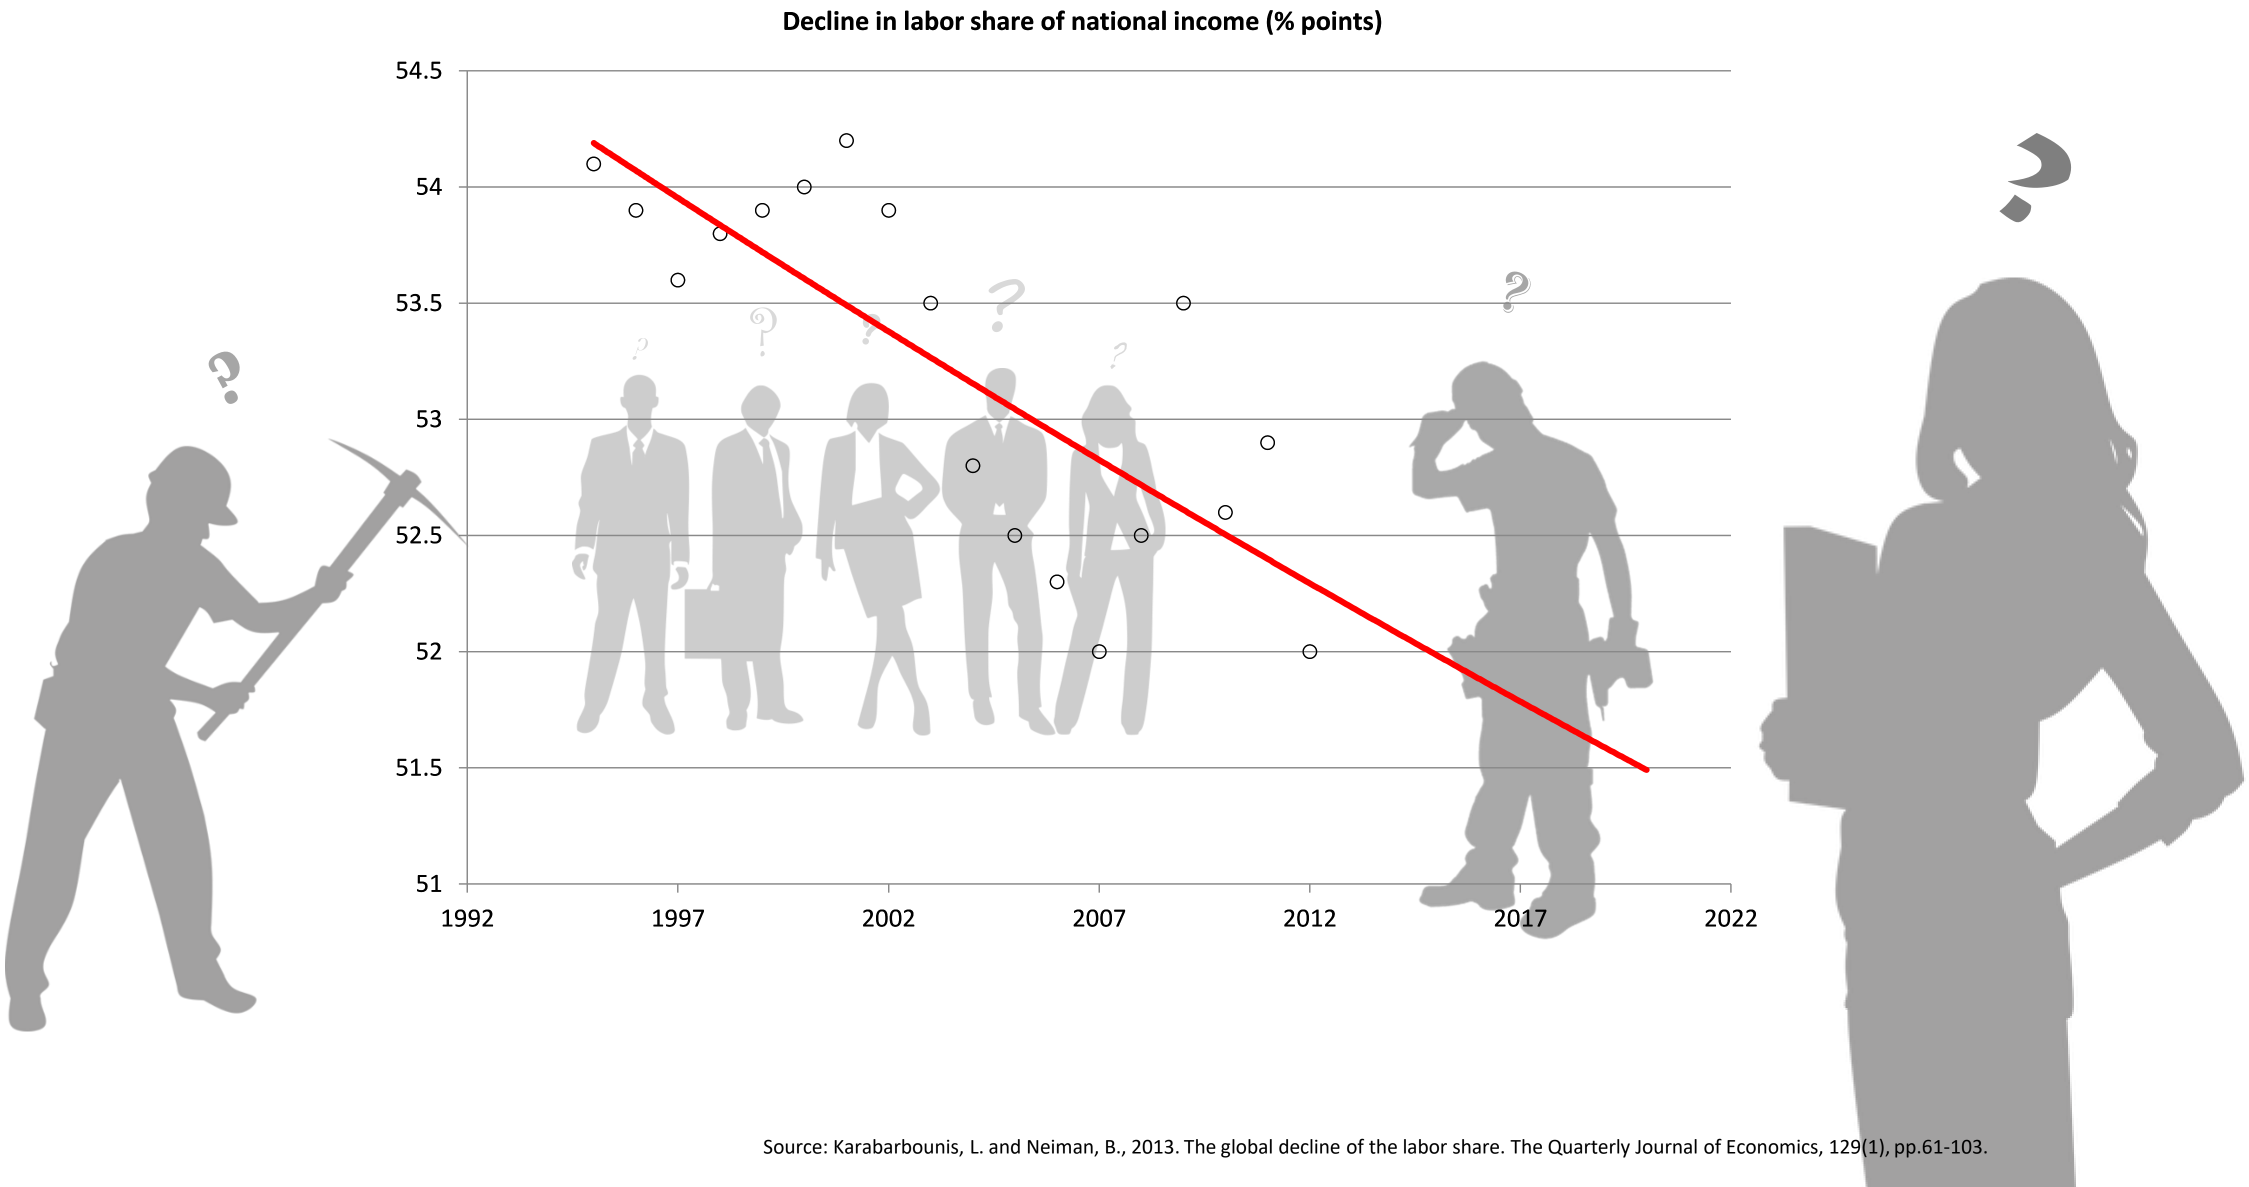

Source: Karabarbounis, L. and Neiman, B., 2013. The global decline of the labor share. The Quarterly Journal of Economics, 129(1), pp.61-103.

ROA

ver.di

zukunft zwei

Maastricht University ROA

IBM

WZB
Berlin Social Science Center

Watson & Future Work

Stockholm University

Tilburg University

ai:conomics

European University Institute

Technequality

ai:nfantry

13

Tallinn University

Future Fit

University of Oxford

Oxford Martin School

Bundesministerium für Arbeits und Soziales

cambridge econometrics

IAB

UM

ver.di

zukunft zwei

Maastricht University ROA

IBM

WZB
Berlin Social Science Center

Watson & Future Work

Stockholm University

Tilburg University

ai:conomics

European University Institute

Technequality

ai:nfantry

13

Tallinn University

Future Fit

University of Oxford

Oxford Martin School

Bundesministerium für Arbeits und Soziales

cambridge econometrics

IAB

Automatisering gaat banen *vernietigen*...
... maar we weten niet *hoeveel*...

1

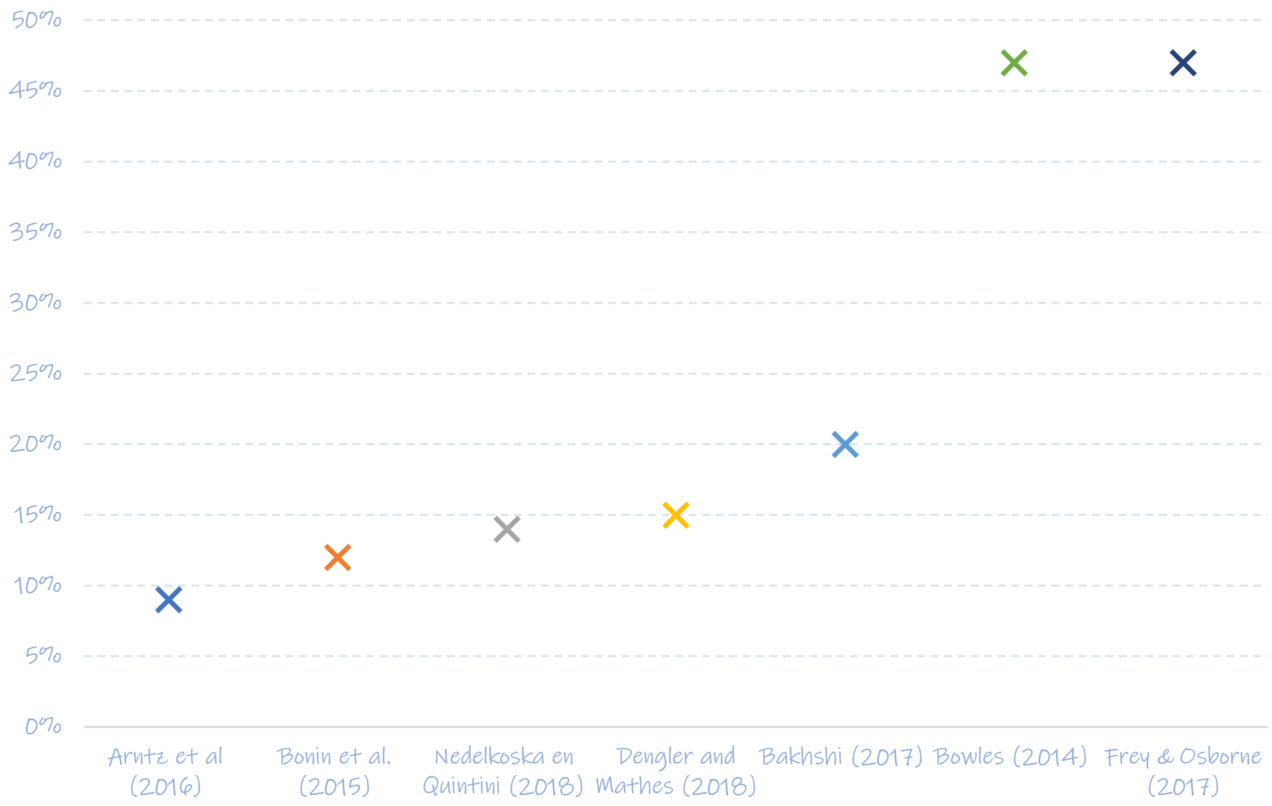

ROA

1

Automatisering gaat banen *vernietigen*…
… maar we weten niet *hoeveel*…

Literatuur

- Alleen focus op deel van de puzzel: tech
- Géén rekening met arbeidsmarktbewegingen
- Rigide assumpties

Onze bijdrage

- Forecasting van knelpunten (POA)
- Redelijke scenario's
- Regio's

ROA

Bronnen: Heald, Smith and Fouarge (2019) - Labour market forecasting scenarios for automation risks: Approach and outcomes, Technequality Report.
Levels, Fregin, Somers (2020). Scenarios for the Impact of Intelligent Automation on Work, Technequality Report

Technequality

1

Automatisering gaat banen vernietigen...

...el...

Employment impact

Figure 3 Employment lost to automation compared to base scenario, by occupation, 2030 (thousands)

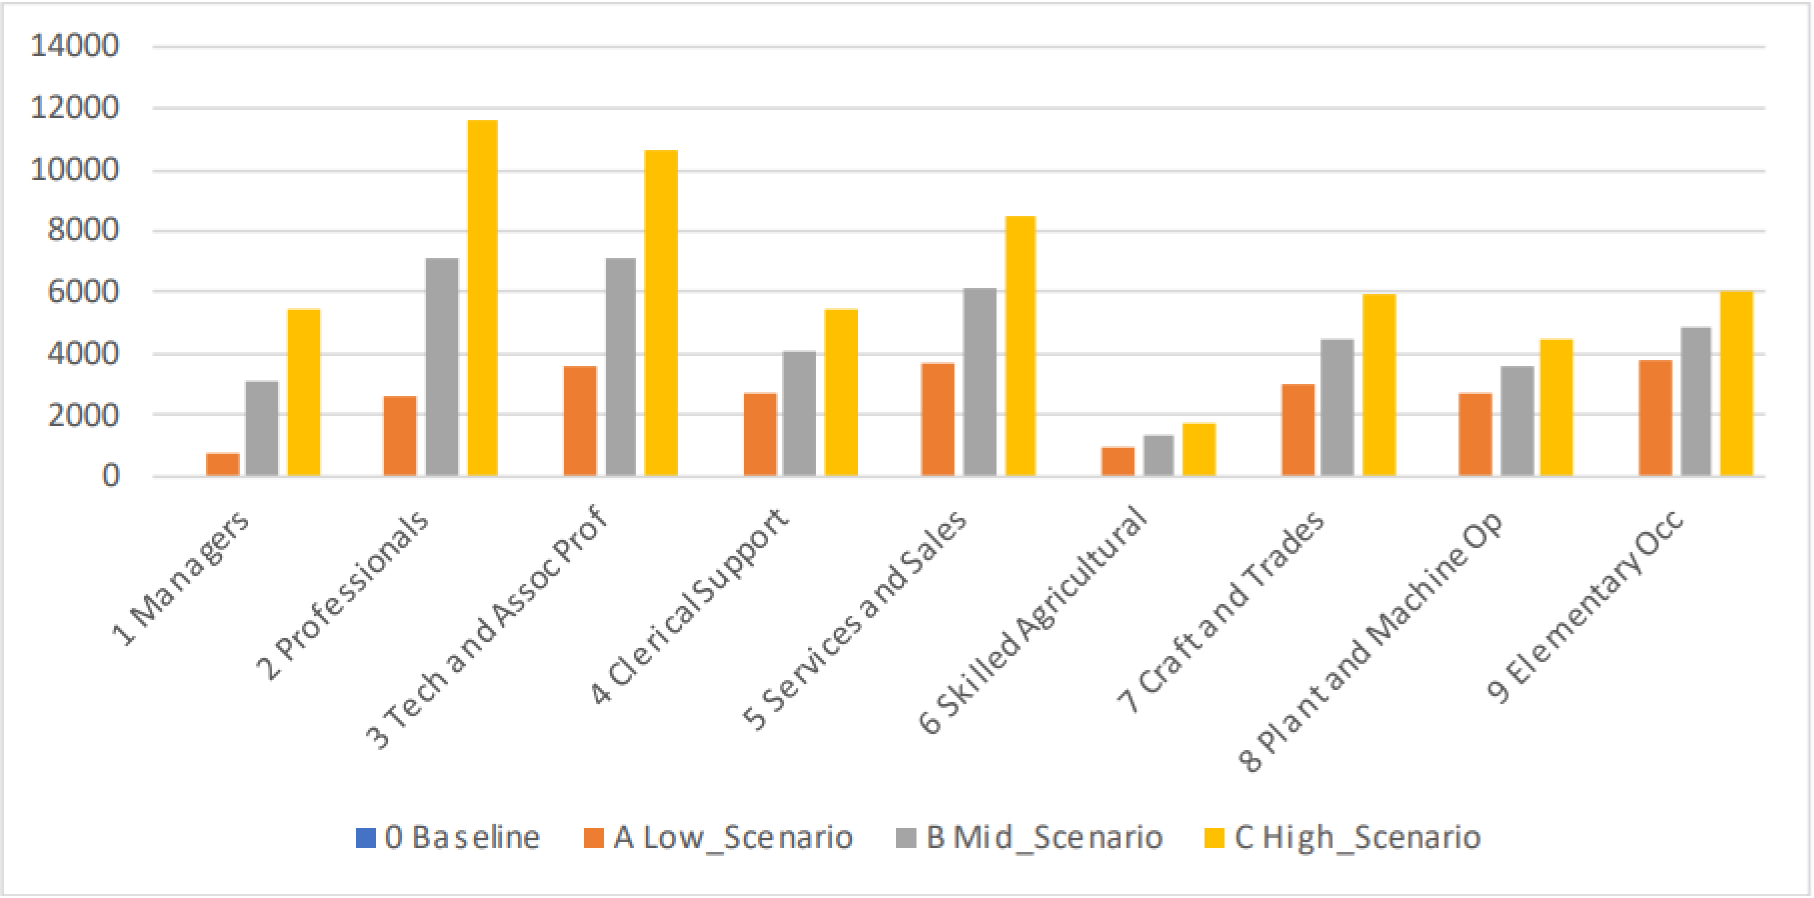

Source: Own calculations based on Cedefop, Eurofound (2018)

and outcomes, Technequality R...

Levels, Fregin, Somers (2020). Sc...

Source: Technequality (2020f)

...port

ROA

Technequality

UM

2a

Automatisering gaat banen creëren...

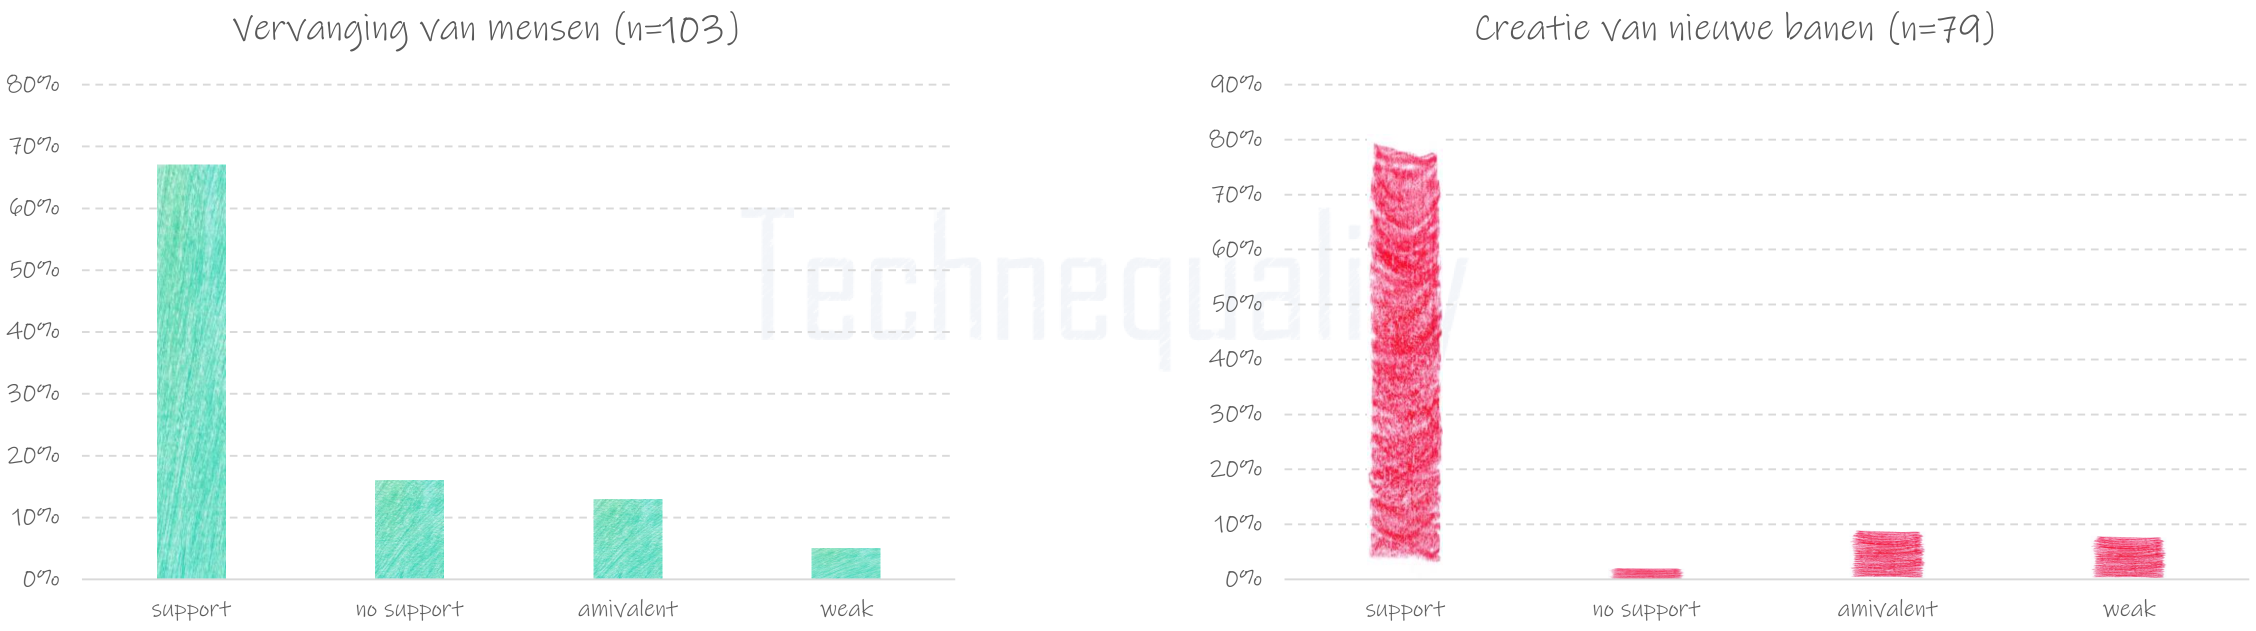

ROA

Source: Kerstin Hötte, Melline Somers, Angelos Theodorakopoulos (2022) Technology and jobs: A systematic literature review. arXiv:2204.01296 [econ.GN]

UM

2b

Automatisering gaat banen creëren...

"About 85 million jobs destroyed, and about 97 million jobs created"

Bron: World Economic Forum (2020) The Future of Jobs Report 2020. Geneva: WEF.

2c

Automatisering gaat banen creëren…
… maar die vragen andere vaardigheden …

Groeiende vraag:
1. Data-analisten en data scientists
2. Ai en machine learning specialisten
3. Big Data specialisten
4. Digitale Marketing- en strategiespecialisten
5. Procesautomatiseringsspecialisten
6. Business developers
7. Digitale transformatiespecialisten
8. Analisten Informatieveiligheid
9. Software en app developers
10. IoT specialisten

Afnemende vraag:
1. Datatypisten
2. Administratieve en secretariële medewerkers
3. Medewerkers accounting, boekhouders, en P&O
4. Accountants en auditors
5. Fabrieksarbeiders
6. Managers in dienstverlening en administratie
7. Klantinformatie- en customer service medewerkers
8. Algemeen en operationeel managers
9. Monteurs en machinereparateurs
10. Opslagmedewerkers

Bron: World Economic Forum (2020) The Future of Jobs Report 2020. Geneva: WEF.

ROA

3

Automatisering gaat banen veranderen ... maar we weten niet *hoe* ...

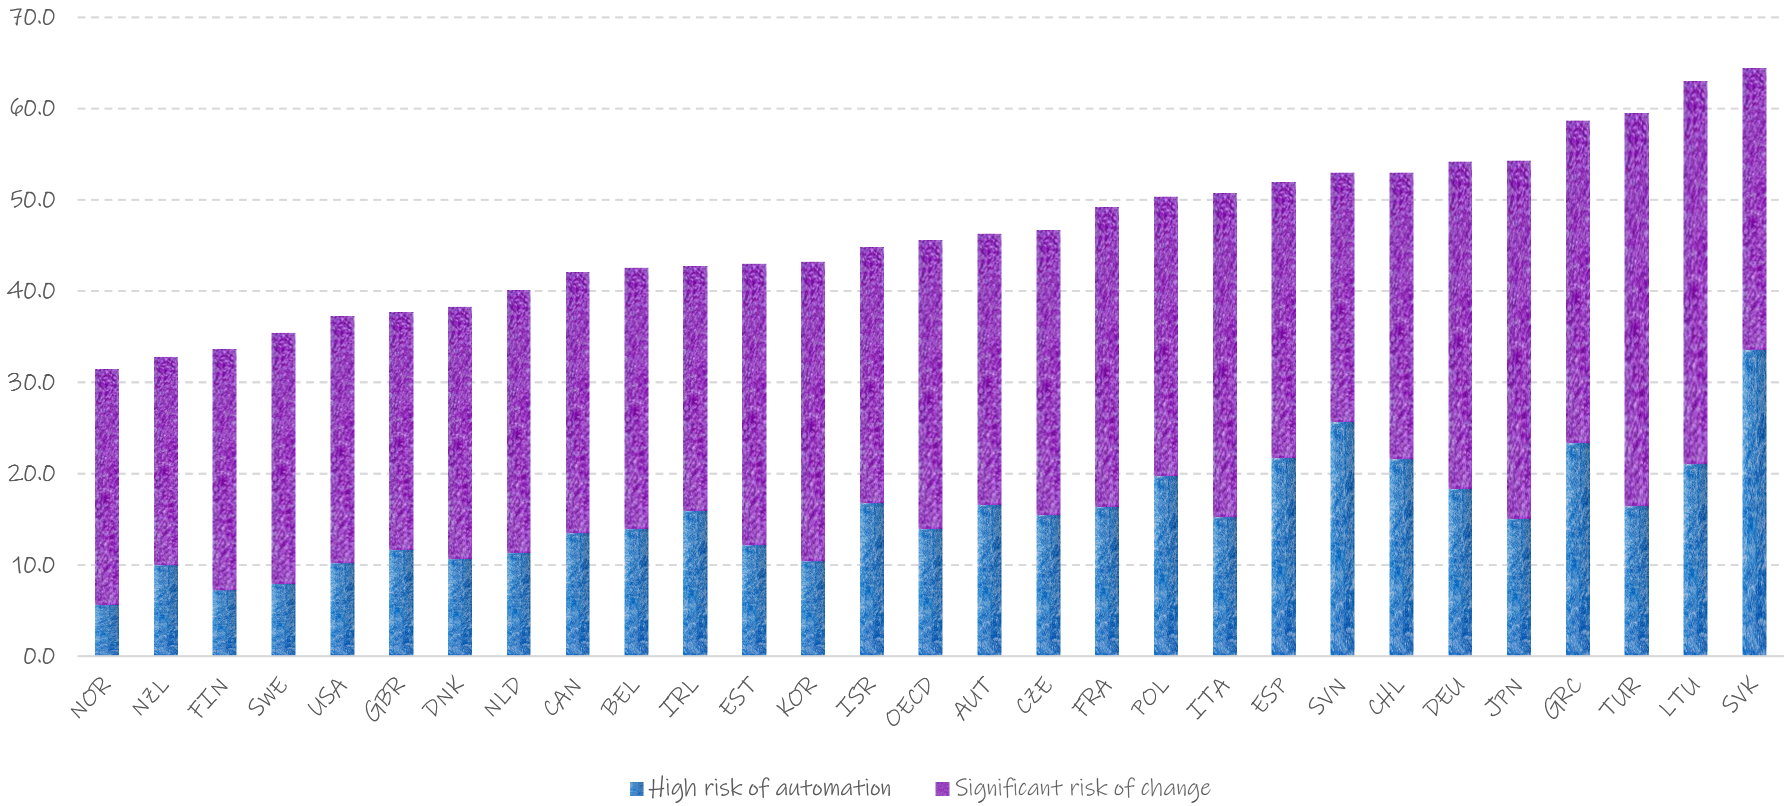

Bron: OECD (2019), OECD Employment Outlook 2019: The Future of Work, OECD Publishing, Paris.
https://doi.org/10.1787/9ee00155-en

4 Automatisering verandert taken
… maar we weten niet *welke*, of *hoe*…

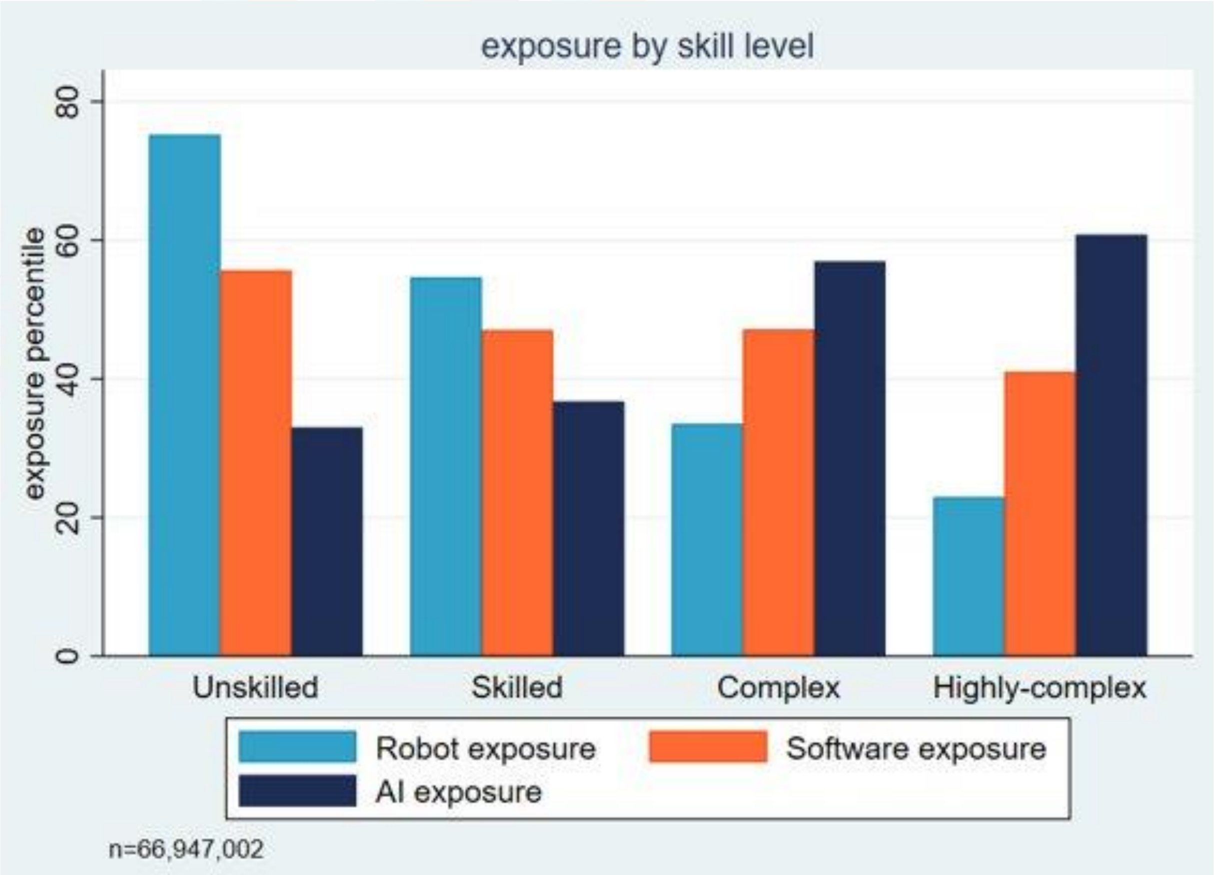

Pelin Ozgül et al.

4 Automatisering verandert taken
… maar we weten niet *welke*, of *hoe*…

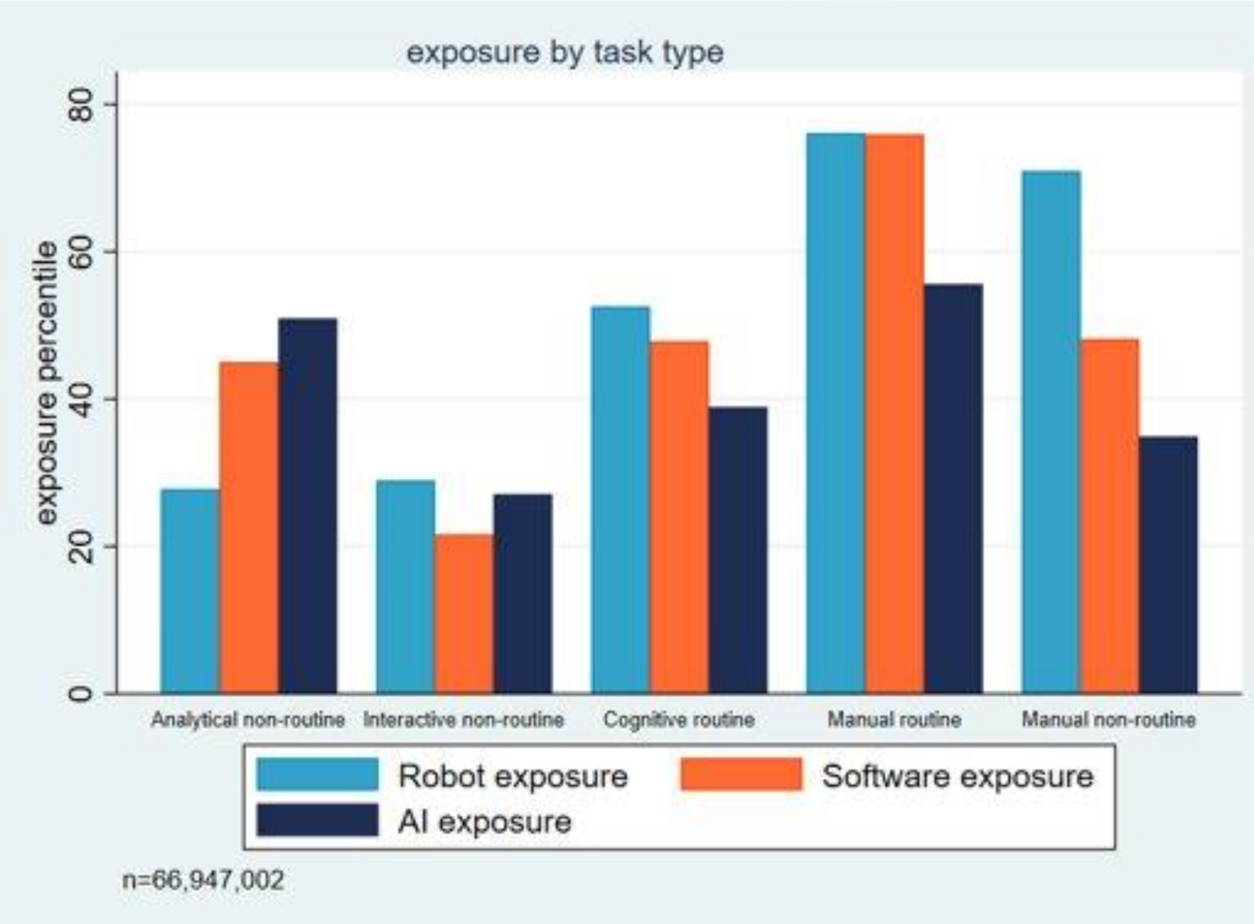

Pelin Ozgül et al.

Menu nrc.nl›

Alice roddelt niet en zal nooit lachen als je je gebit uitdoet

Maatje Zorgrobot Alice is nog lang geen volwaardige gesprekskameraad. Maar van eenzame ouderen hoeft ze niet per se menselijk te zijn.

Liza van Lonkhuyzen 18 oktober 2019 Leestijd 5 minuten

Arbeidsmarkt van de toekomst

Ándere transities:

1. Automatisering
2. Energietransitie
3. Vergrijzing/ ontgroening
4. (de-)globalisering
5. Geostrategische industriepolitiek

Vraag naar skills

Arbeidsmarkt van de toekomst

Ándere transities:

1. Automatisering
2. Energietransitie
3. Vergrijzing/ ontgroening
4. (de-)globalisering
5. Geostrategische industriepolitiek

Vraag naar skills

ROA

Maastricht University ROA

Rapport: “Inzicht in arbeidsmarktknelpunten voor de uitvoering van het klimaatbeleid”

Planbureau voor de Leefomgeving

1 De energietransitie verandert de vraag naar arbeid…
… maar hoe hangt af van klimaatbeleid …

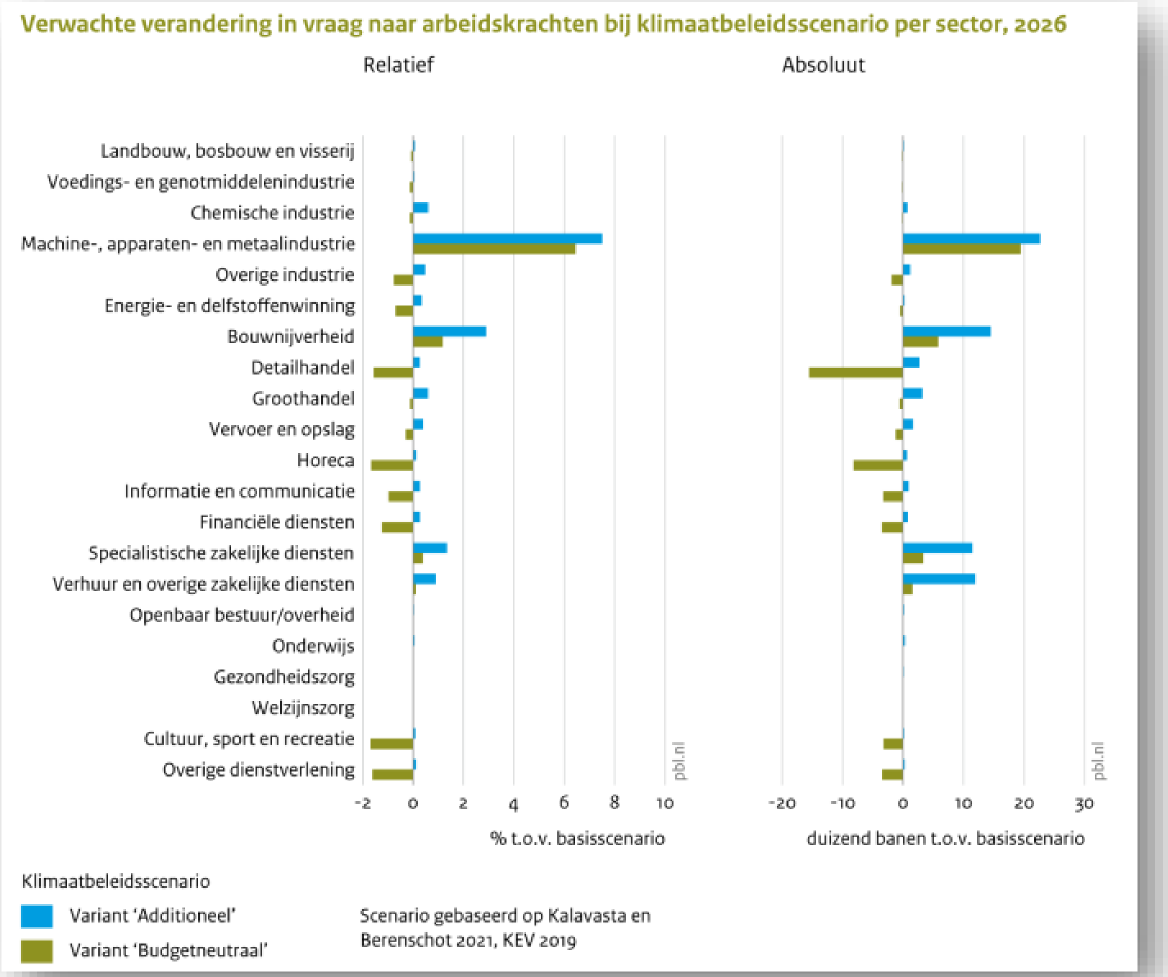

2

Tekort aan geschoold personeel vormt een

zeer belangrijke hindernis

bij het halen van de klimaatdoelen...

3 Grote regionale spreiding, óók door beleid …

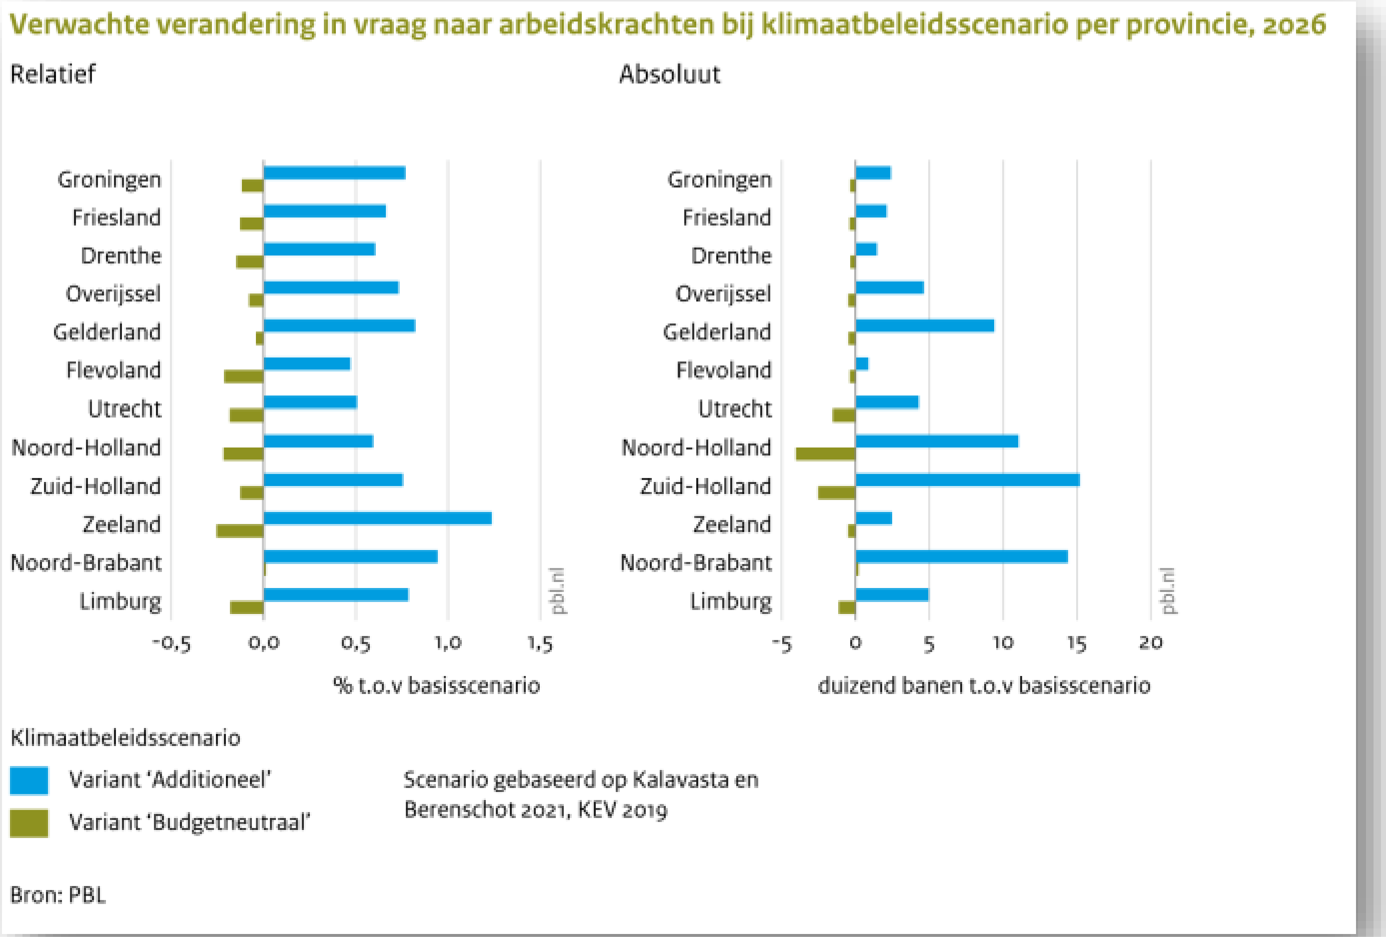

Conclusies?

Conclusies tot zo ver…

Transities zullen de arbeidsmarkt *ondoelmatiger* maken …

… en dat kán productiviteitsgroei verhinderen …

… leiden tot toename werkloosheid …

… en grotere maatschappelijke ongelijkheid …

… en verhindert het halen van cruciale beleidsdoelen …

… en we weten heel veel nog niet …

… wel: impact op beroepenstructuur én beroepen …

Wat is nodig?

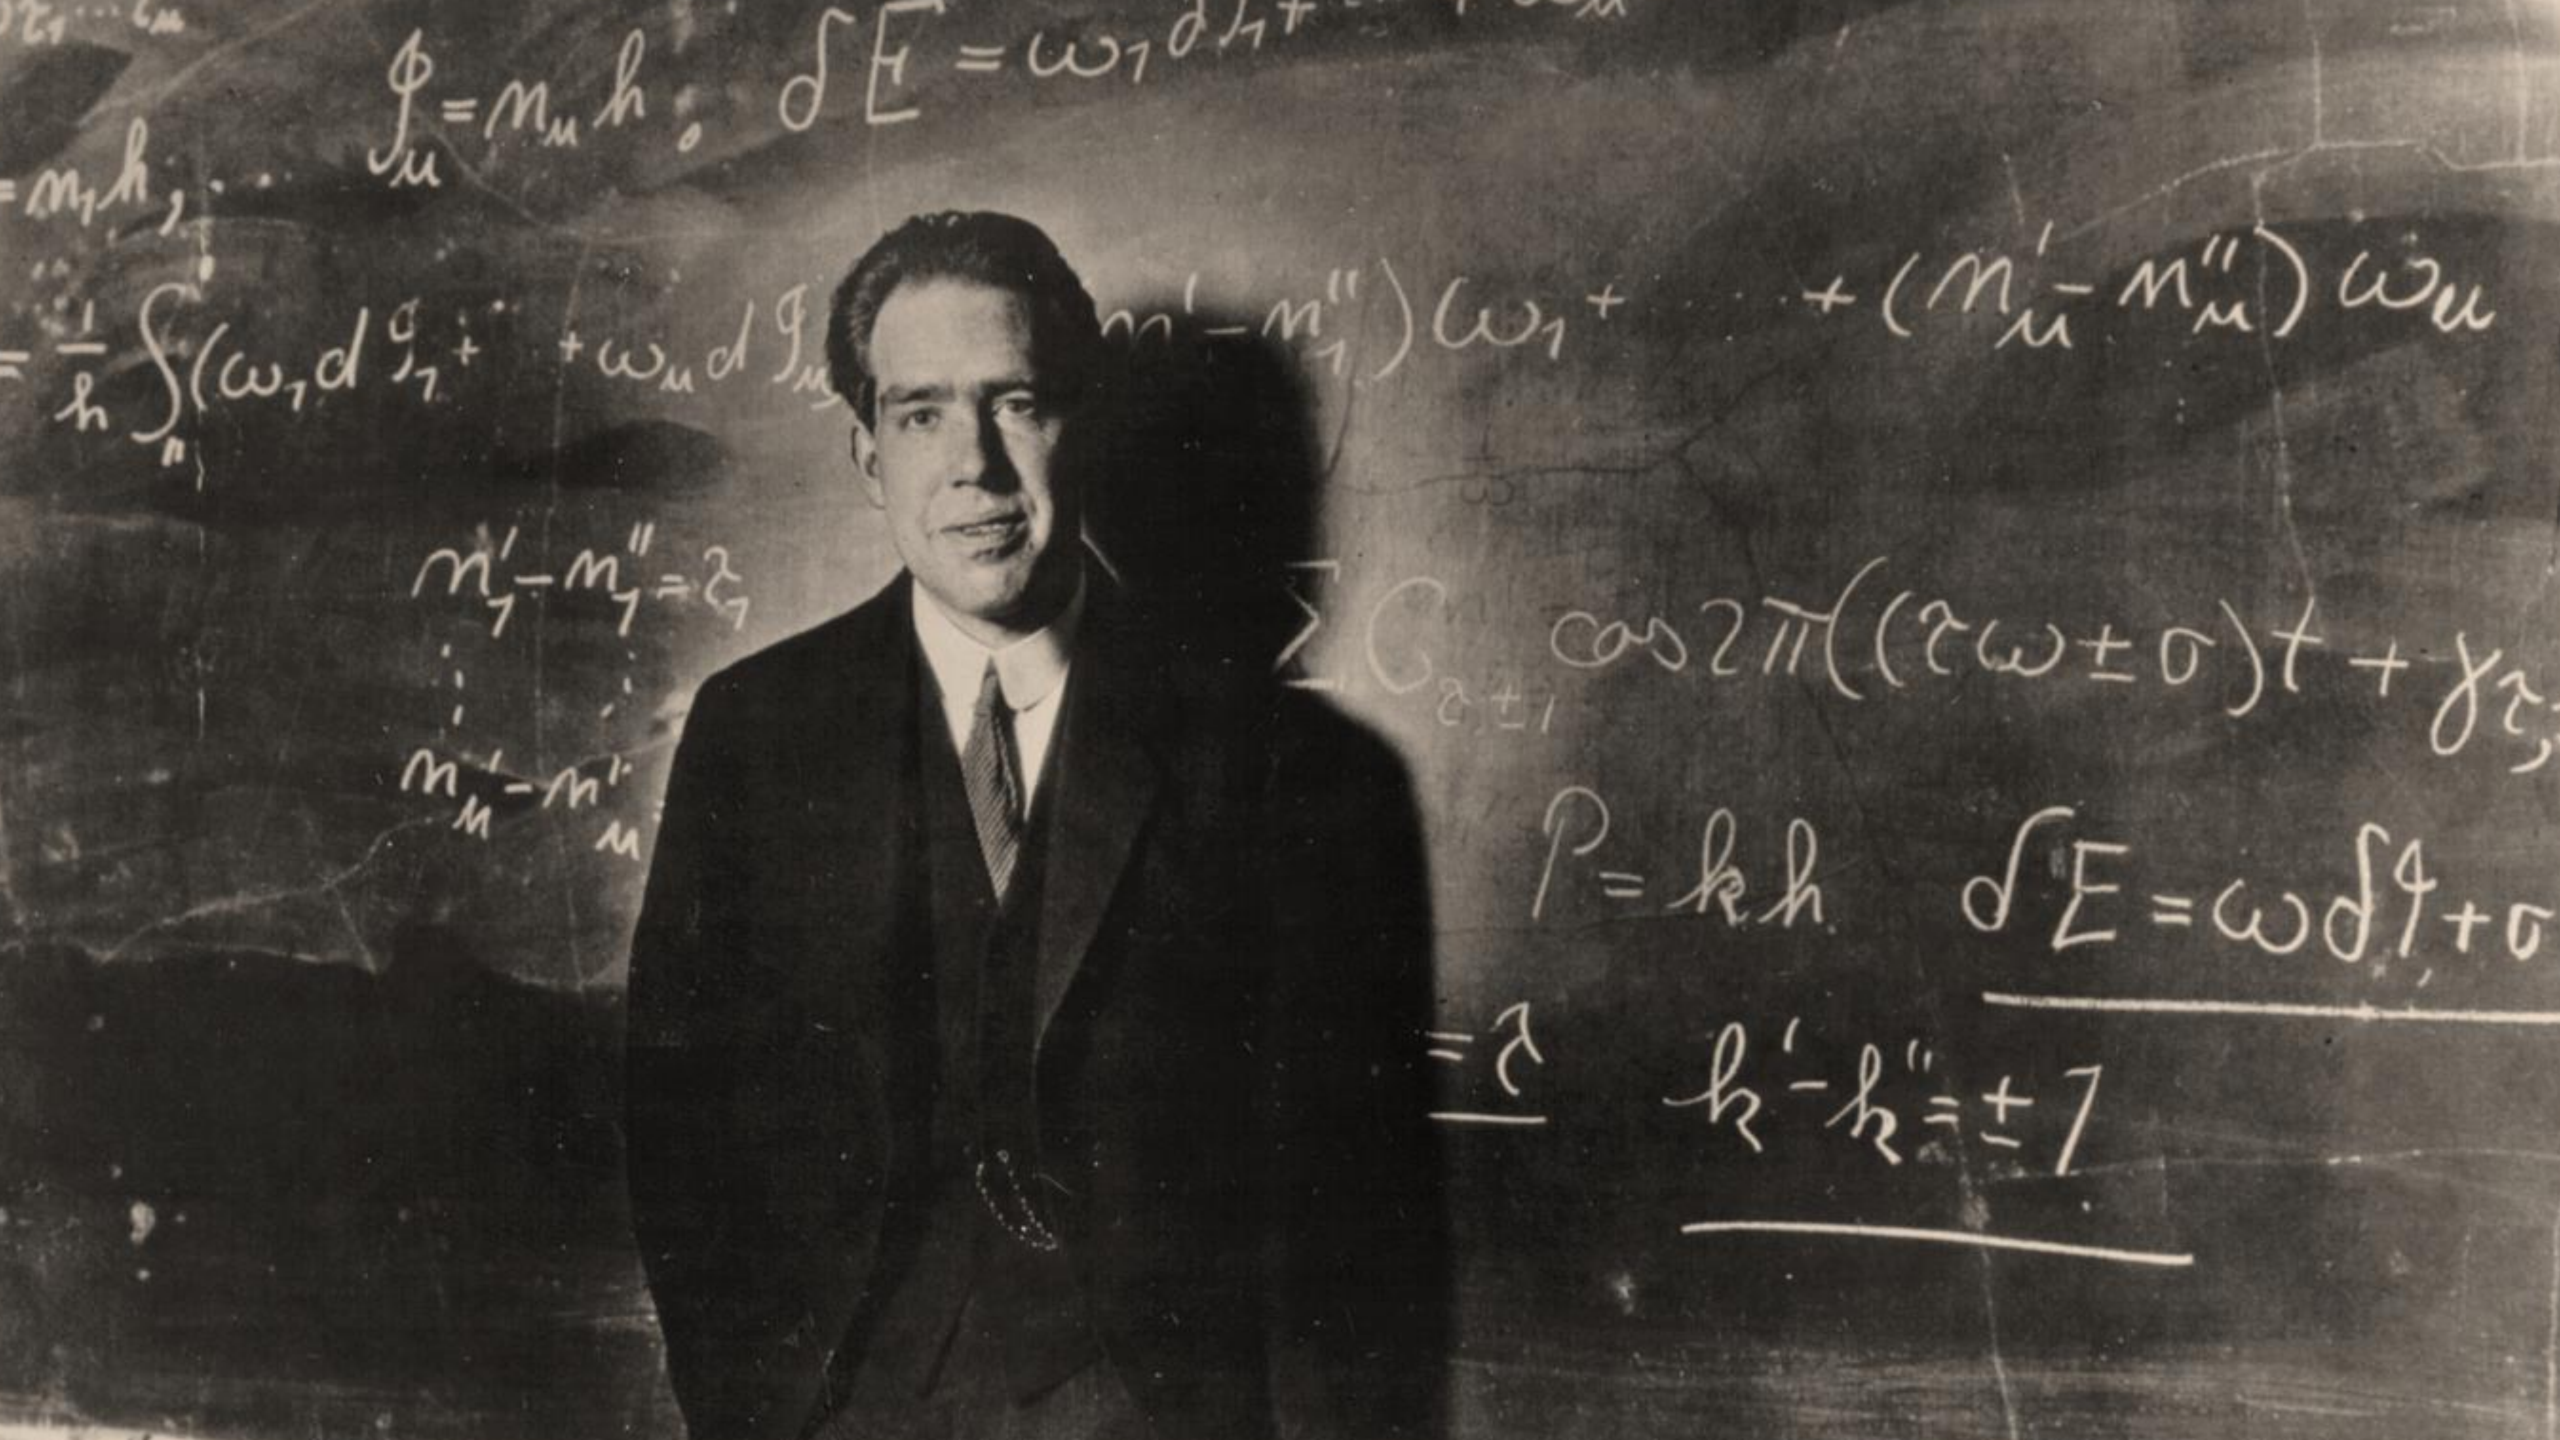

Nodig…

Innovaties op voorspellend model…

… van onderwijs naar skills …

… van beroepen naar taken …

… betere regionale voorspellingen …

… Inzicht in veranderende skills vraag …

… door onderzoek naar beroepen …

… goed begrip aanbod van skills …

… goed begrip vraag naar LLO …

ROA

REALLY???

LLO als oplossing ?!

1. Heel veel mensen nemen niet deel aan cursussen!

Vooral werkenden leren bij...

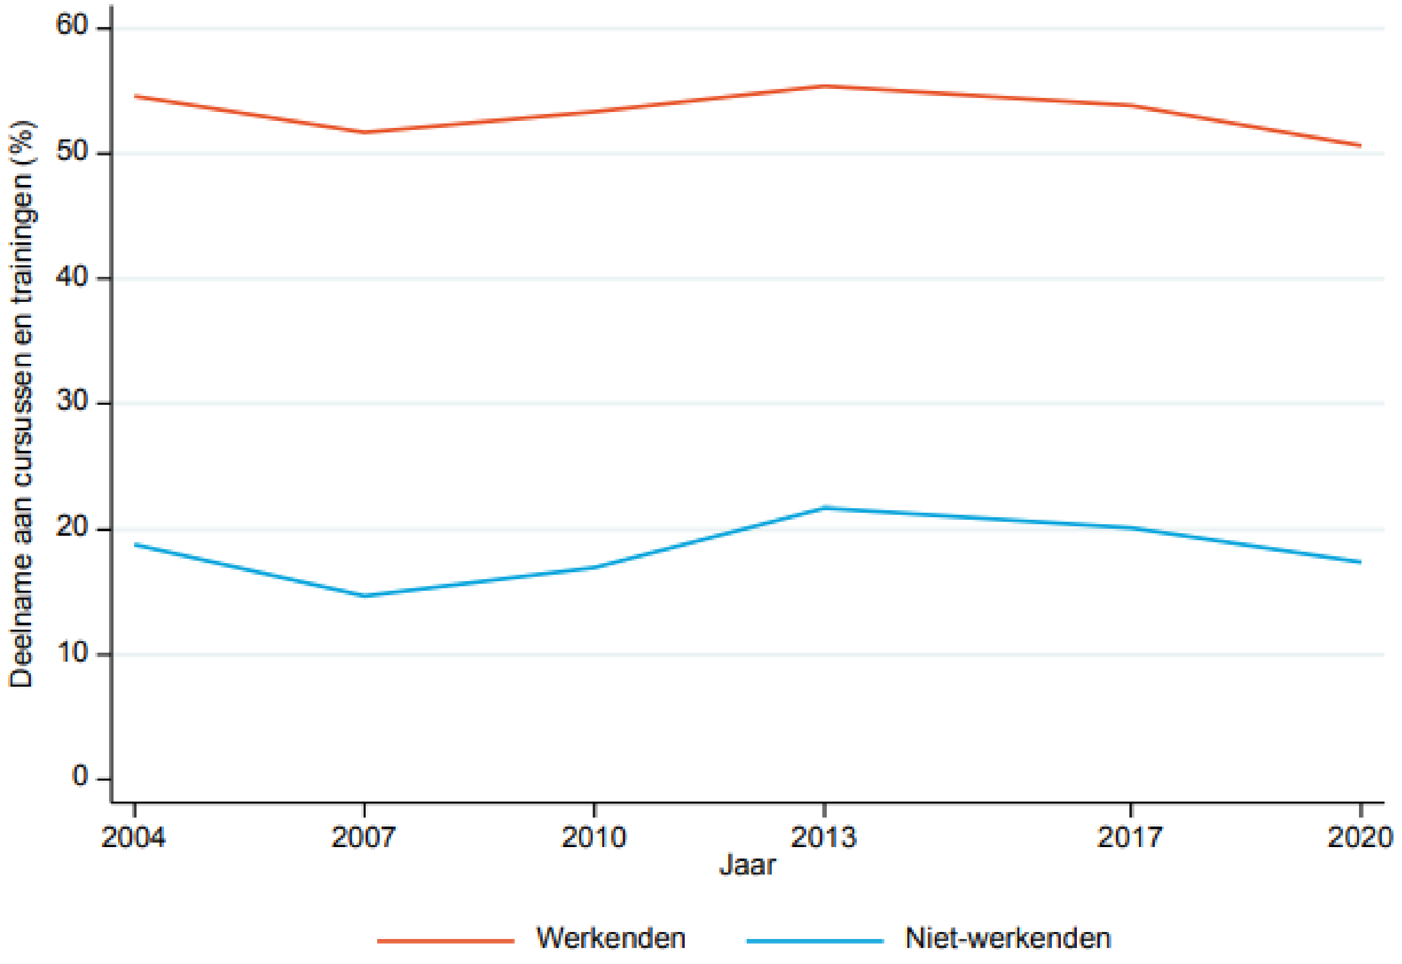

Künn, A., Abbink, H., Baumann, S., van Elferen, S., & Fouarge, D. (2022). Leven lang ontwikkelen in Nederland. ROA-R-2022/1 https://doi.org/10.26481/umarep.2022001

Maar ouderen nemen minder vaak deel …

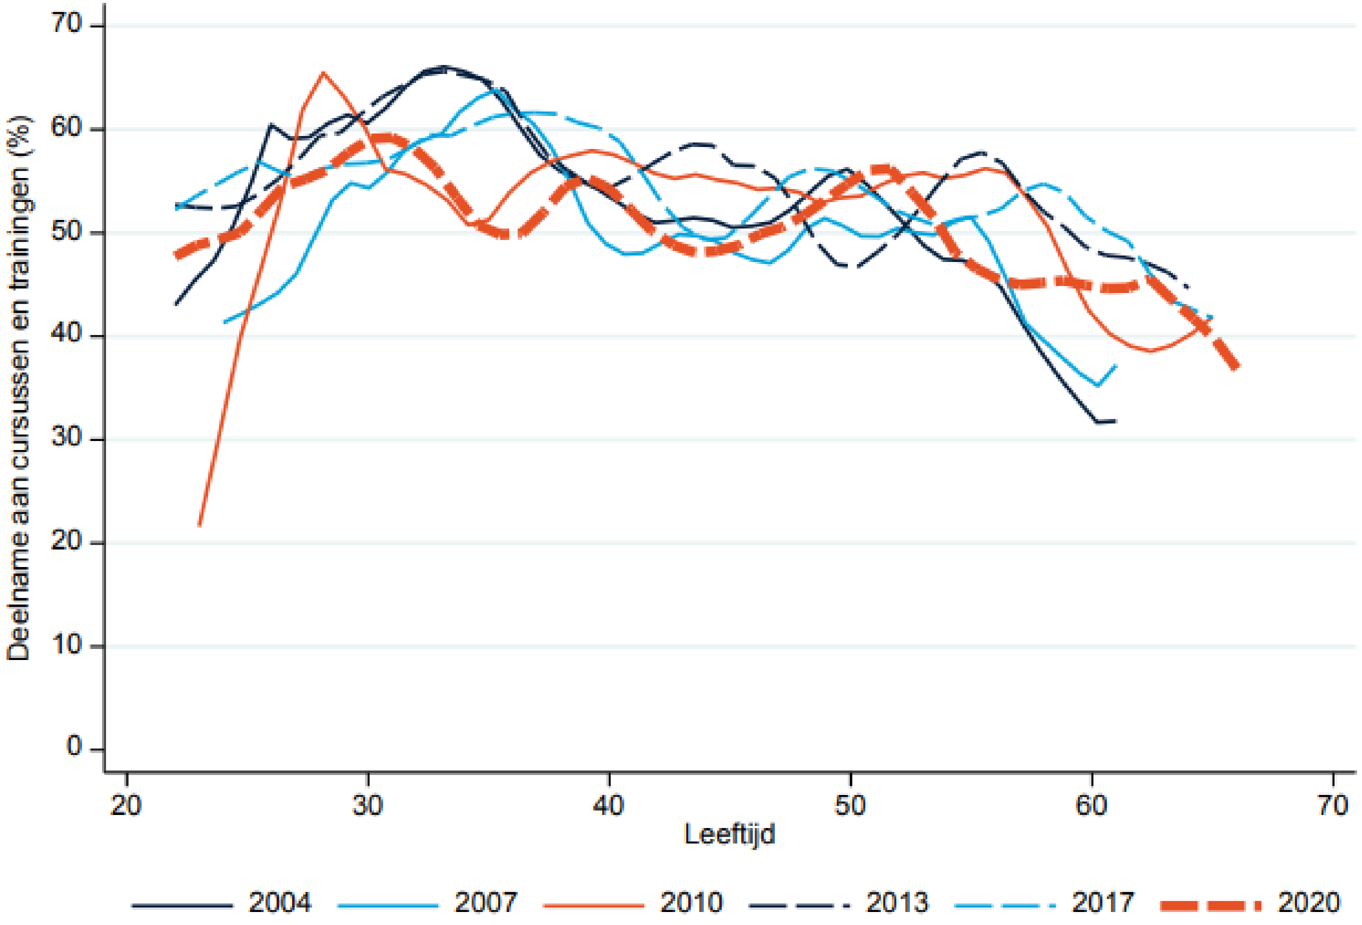

Künn, A., Abbink, H., Baumann, S., van Elferen, S., & Fouarge, D. (2022). Leven lang ontwikkelen in Nederland. ROA-R-2022/1 https://doi.org/10.26481/umarep.2022001

… en leren gebeurt *minder* in sterker automatiseerbare beroepen …

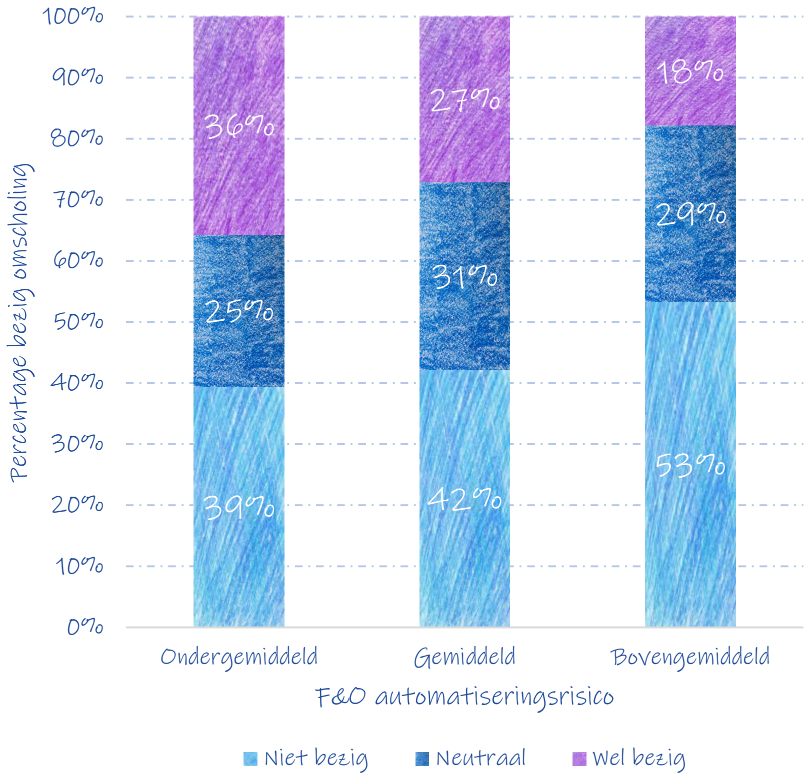

Bron: Jansen, G., Janssen, S., Künn, A., Levels, M., & Voermans, L. (2021). Bij- en omscholing en automatisering in Nederland: Kansen en beperkingen. Universiteit Twente. ROA External Reports.

… ook lager opgeleiden trainen veel minder…

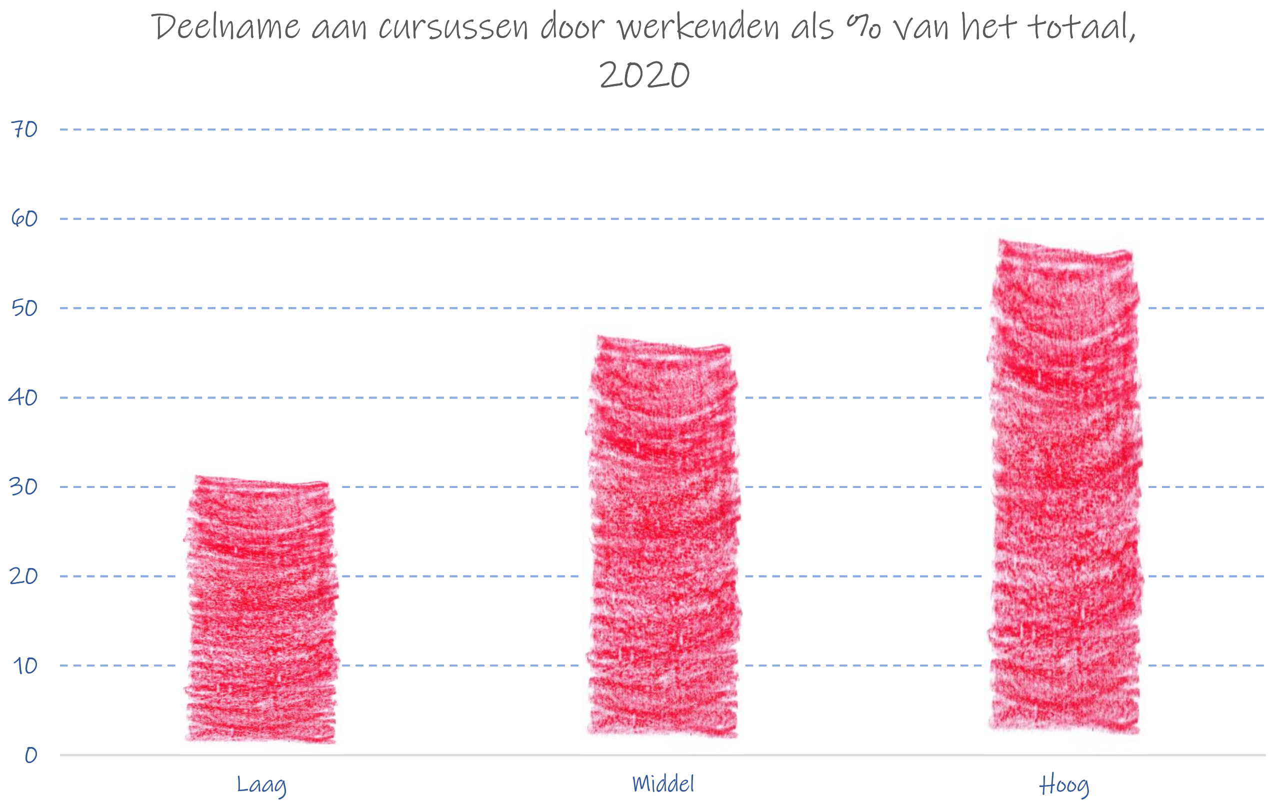

Bron: ROA LLL Enquête 2004-2020, uit De Grip (2021) Van een leven lang leren naar een leven lang ontwikkelen: De veranderende betekenis van het post-initieel leren.. Afscheidsrede als hoogleraar aan Maastricht University.

… evenals flex wekers …

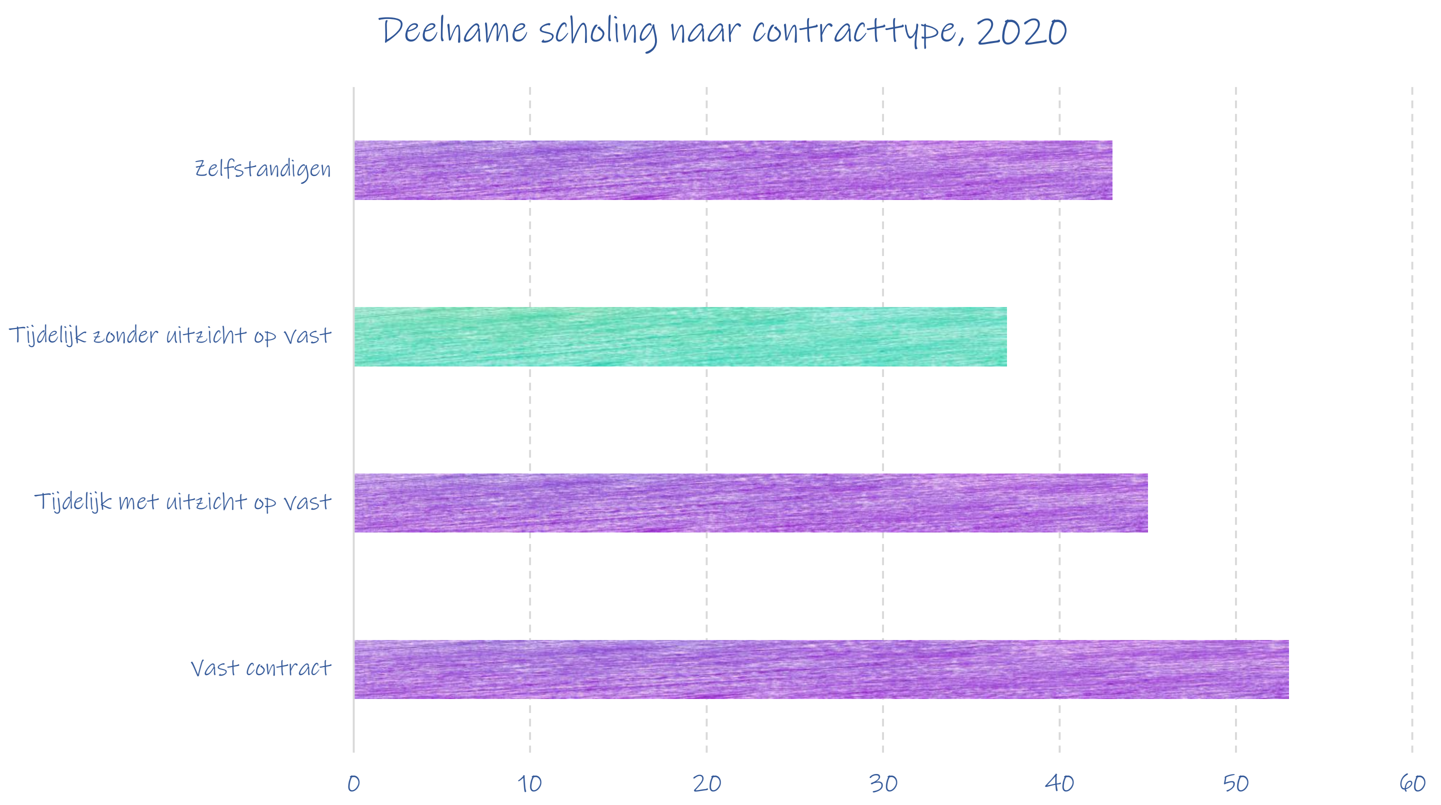

Bron: ROA LLL Enquête 2004-2020, uit De Grip (2021) Van een leven lang leren naar een leven lang ontwikkelen: De veranderende betekenis van het post-initieel leren.. Afscheidsrede als hoogleraar aan Maastricht University.

LLO als oplossing ?!

2. LLO leidt niet altijd tot betere uitkomsten !

... Lonen ...

... Productiviteit ...

... Re-integratie ...

... Gepercipieerde inzetbaarheid ...

... Carrière perspectieven ...

Dan maar geen LLO?

NEE !

Inzicht leidt tot oplossingen !

Waarom neemt niet iedereen deel aan LLO-aanbod?

... Werknemers in beroepen met een hoog automatiseringsrisico ...

- Zien de noodzaak niet

ROA

Want: meerendeel werknemers denkt dat beroep niet automatiseerbaar is

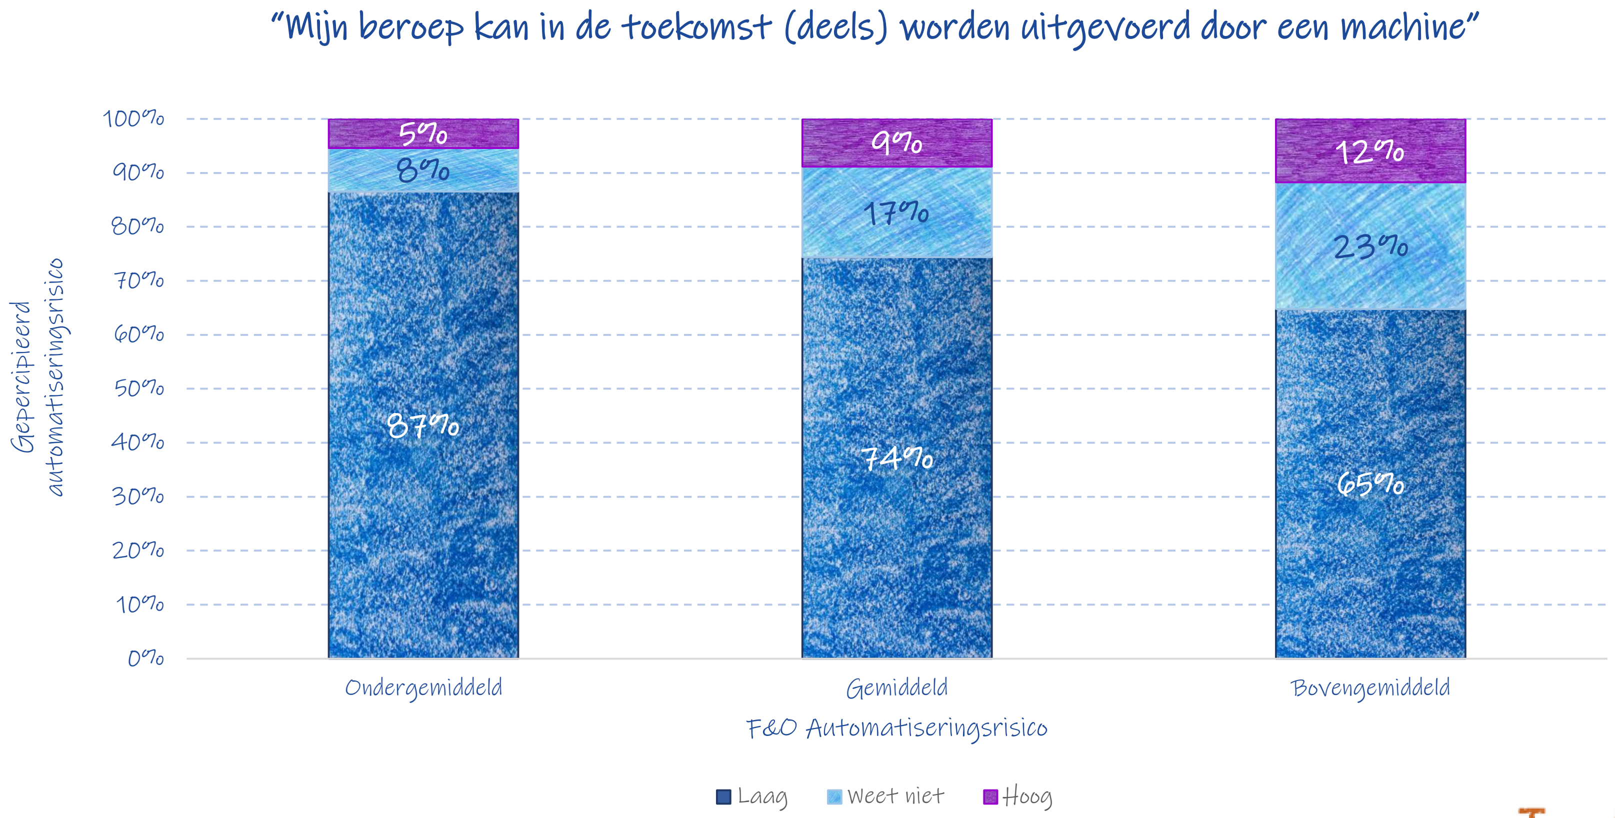

Bron: Jansen, G., Janssen, S., Künn, A., Levels, M., & Voermans, L. (2021). Bij- en omscholing en automatisering in Nederland: Kansen en beperkingen. Universiteit Twente. ROA External Reports.

Inzicht leidt tot oplossingen !

Waarom neemt niet iedereen deel aan LLO-aanbod?

… werknemers in beroepen met een hoog automatiseringsrisico …

- Zien de noodzaak niet
- *Werkgevers minder bereid* *(Heß et al. 2023; Künn et al., work in progress)*
- *Wat moet er geleerd worden?*

ROA

Inzicht leidt tot oplossingen !

Waarom neemt niet iedereen deel aan LLO-aanbod?

... Oudere werknemers...

- Scholing sluit niet goed aan

... Praktisch opgeleiden ...

- Lage self-efficacy
- *Zien noodzaak niet*
- *Wat?*
- *Waar?*

Inzicht leidt tot oplossingen !

Waarom neemt niet iedereen deel aan LLO-aanbod?

… Werkzoekenden …

- Financiële barrière
- Welke skills ?

… Flex workers …

- Werkgevers minder bereid
- *Zien noodzaak niet*

Inzicht leidt tot oplossingen !

Waarom neemt niet iedereen deel aan LLO-aanbod?

... Werkenden in MKB ...

- Minder mogelijkheden

ROA

Evidence based oplossingen !

Vouchers

- Beperkte houdbaarheidsduur vergroot gebruik (Grip)
- Hoogste rendement indien keuze voor geld of tijd (Fleuren etal. 2020)
- Vouchers werken voor laagopgeleiden (Görlitz 2010) en in MKB (Hidalgo et al. 2014)

Let op! Vaak is geld niet het issue, en als, dan 100% vergoeding nodig (Görlitz & Tamm 2017)

Evidence based oplossingen !

Informatie interventies

- Werknemers informeren over de automatiseringskansen van hun beroep vergroot trainingsintentie (Lergetporer et al. 2023)
- *Algemene informatie over automatiseringsrisico verhoogd trainingsdeelname onder werkzoekenden* *(Leduc & Tojerow, work in progress)*
- *Hoe informeer je flex workers het beste?* *(Wittich, Kunn, Montizaan, work in progress)*

Evidence based oplossingen !

Training design

- Training buiten werktijd reduceert bereidheid (Künn et al. 2018)
- *Online cursussen niet in trek* (Künn et al. 2018)
- *Volgen van CVET tijdens werktijd verhoogt deelname* (Rueter 2022)
- *Praktijkcompenent en onderwijsvorm* (Künn et al., work in progress)

Inzicht leidt tot oplossingen !

Waarom draagt LLO niet bij aan betere uitkomsten?

- 30% van de trainingen kunnen niet ingezet worden op het werk (Van Eldert,"Fouarge & Künn 2018)
- *Ouderen hebben baat bij andere trainingsvormen en –inhoud* (Zwick 2015)
- *Veel mensen kiezen training o.b.v. voorkeuren of content waar ze al goed in zijn* (Borghans)

ROA

Evidence based oplossingen !

Training design

- Cursusdesign moet aangepast worden op preferenties en randvoorwaarden
- *We hebben beter zich nodig op de vraag naar vaardigheden*

Als alles samenkomt werkt LLO wel!

Als alles samenkomt werkt LLO wel !

- LLO kan (bedrijfs)productiviteit vergroten, ook van collega's (De Grip & Sauermann 2012)
- *LLO kan gepercipieerde inzetbaarheid vergroten (Van Eldert, Fouarge & Künn 2018)*
- *LLO kan gerelateerd zijn aan*
  - *Hogere lonen* (Barrett & O'Connell 2001)
  - *Meer baantevredenheid*
  - ...

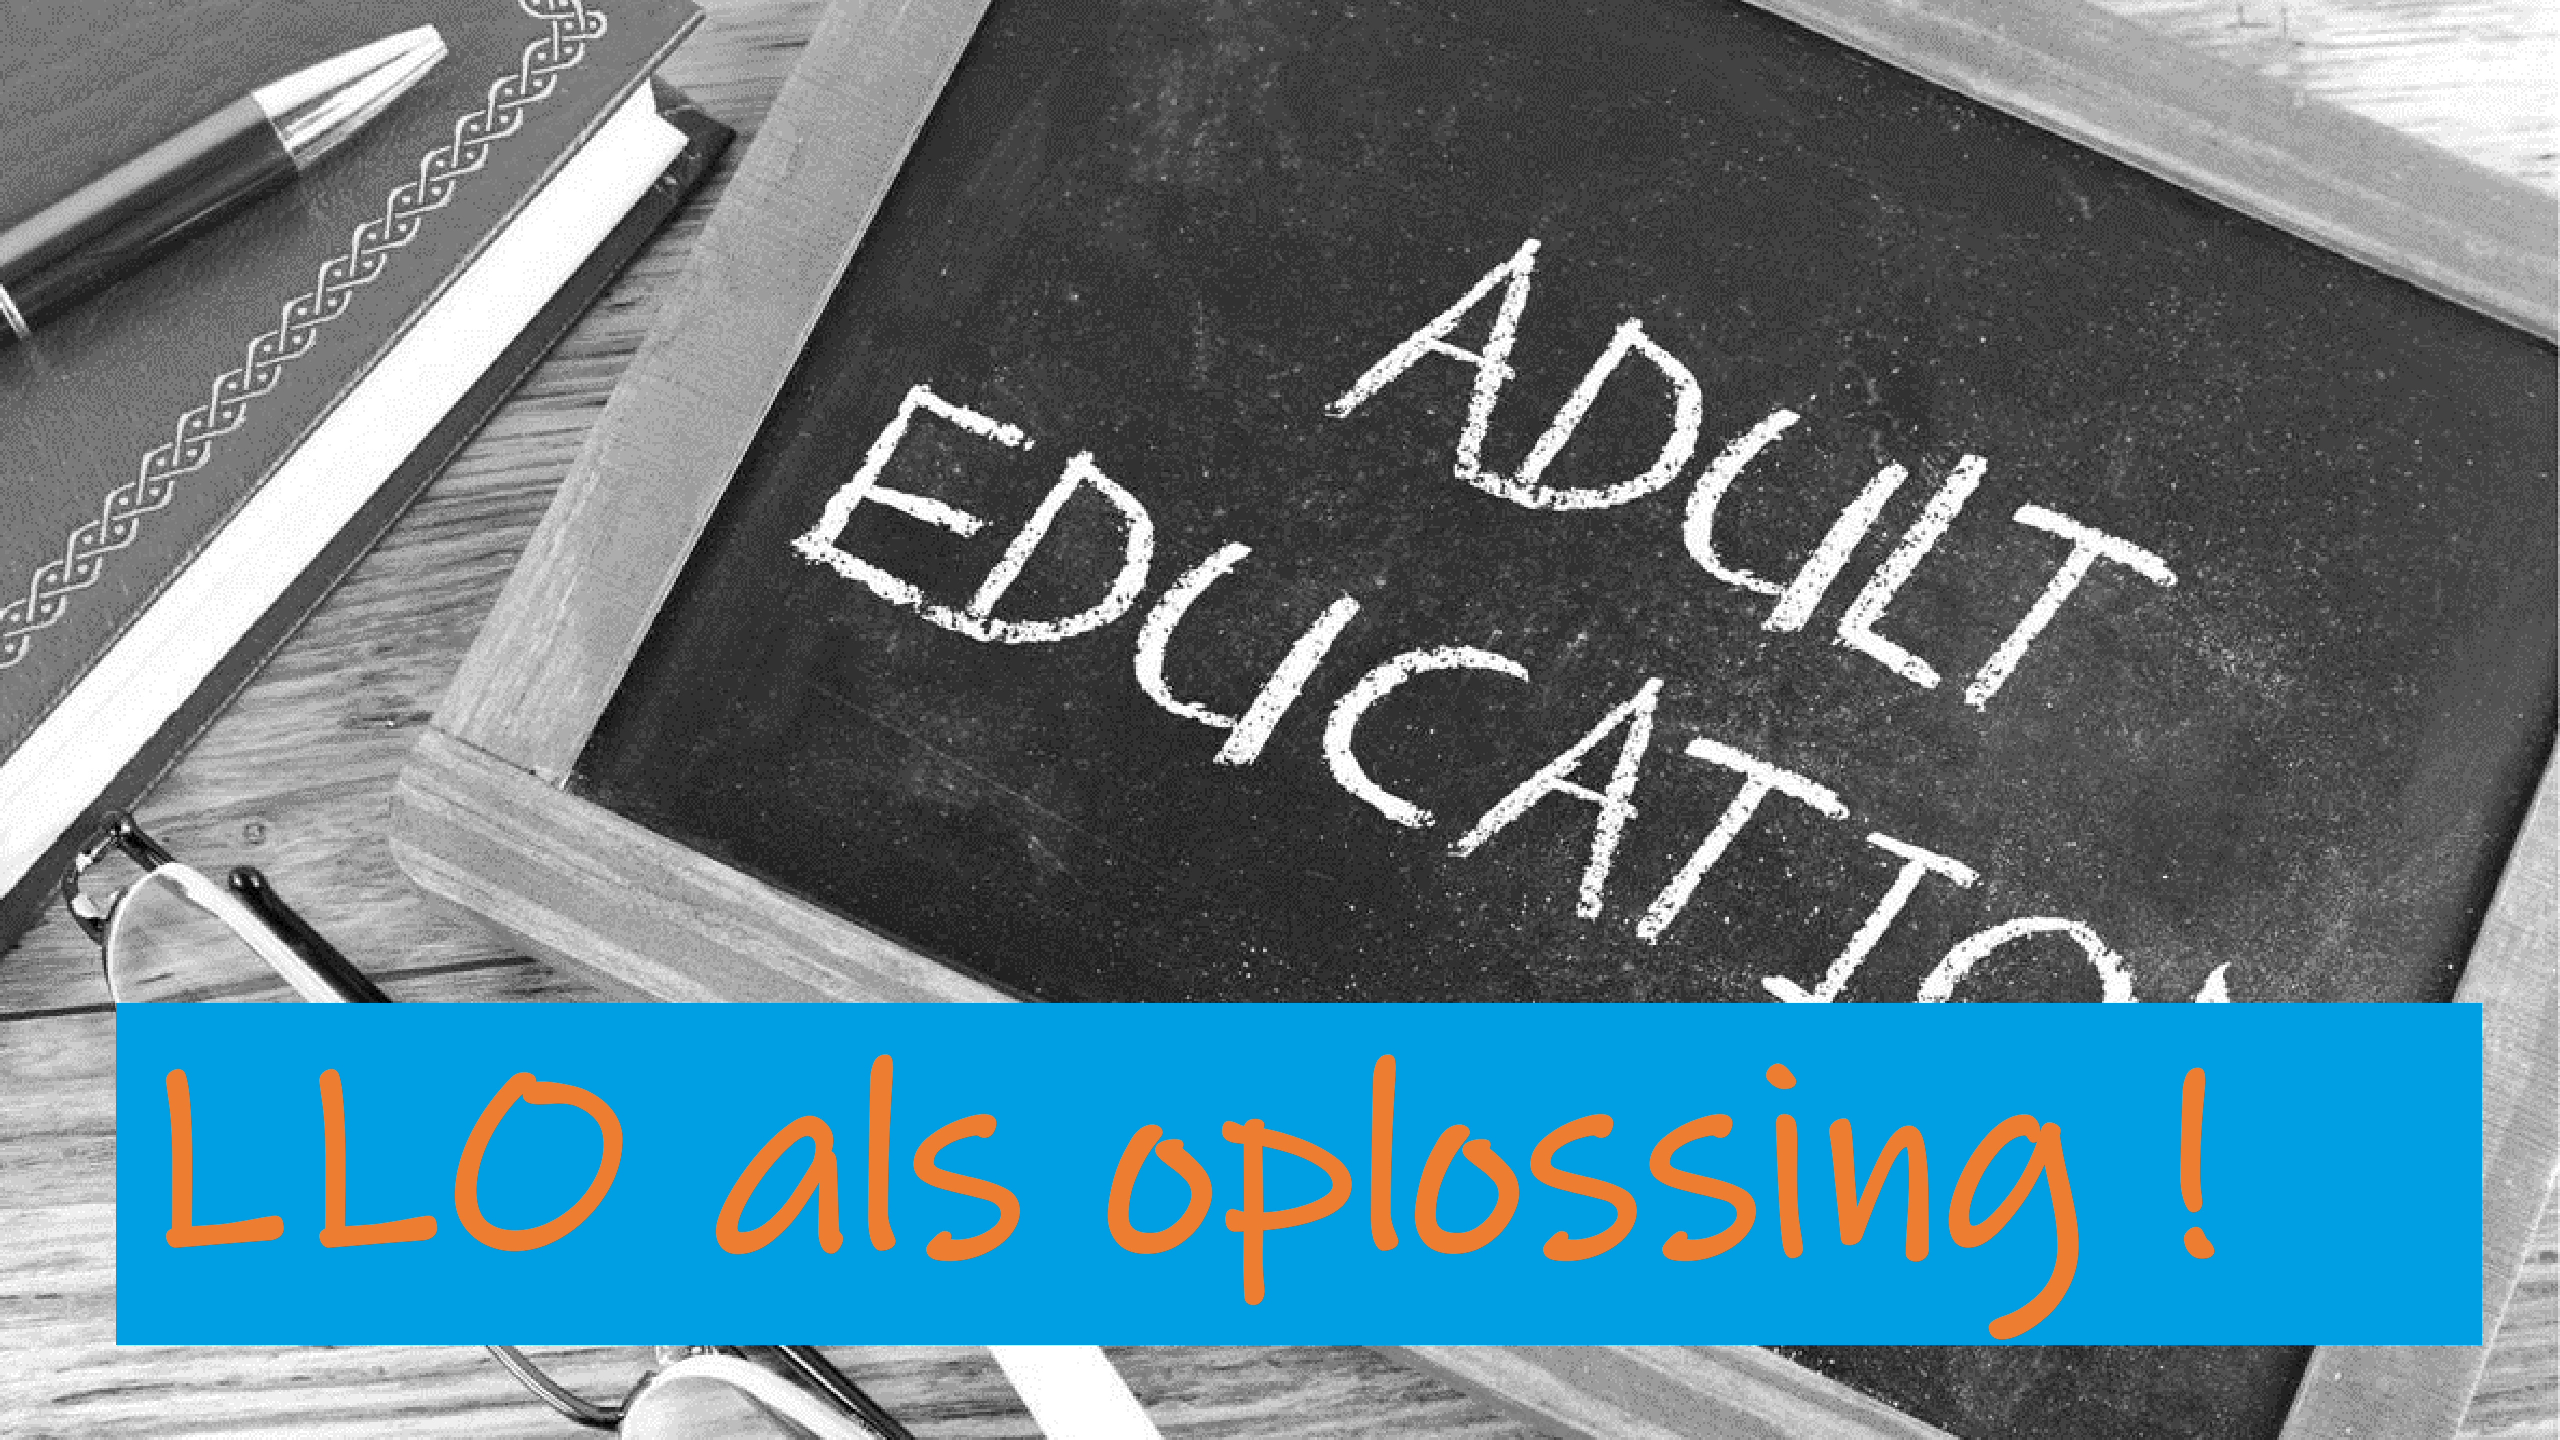

LLO als oplossing !

LLO als oplossing !

Indien:

- Beter zicht op vraag naar vaardigheden
- LLO beter ingebed is op de werkvloer
- Beter zicht op barrières en stimulerende aspecten
- Gerichte interventies

- LLO aanbod moet aangepast worden aan de uitdagingen!

Thx