Economic Possibilities for Our Children: AI and the Future of Work, Education, and Leisure

Miles Brundage
Consortium for Science, Policy, and Outcomes
Virtual Institute of Responsible Innovation
Arizona State University

Overview

- Historical Context
- Goals/Designs in AI Research
- Applications and their Impacts
- Work
- Education
- Leisure
- Basic Income
- AI Progress
- Engaging the Public about AI
- Next Steps

Economic Possibilities for Our Grandchildren

- Written 85 years ago
- Expected six-fold rise in US per capita GDP
- Expected decline in working hours, rise in leisure by 2030.
- Skidelsky and Skidelsky (2012) on what he got wrong:
  - Wants vs. needs
  - Pressure to work
  - Enjoyment of work

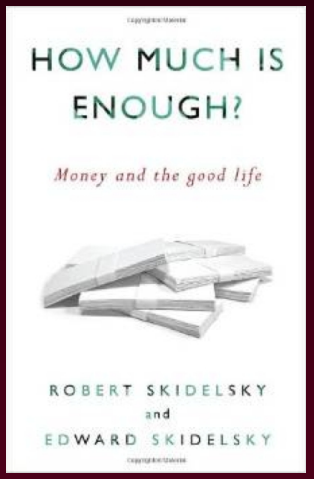

What's Changed Since 1930?

- Progress in AI/robotics/tech in general
- Types of jobs available
- More/better education
- Wider workforce participation

What's Changed Since 1930 (cont'd)?

- More people involved in global economy
- Roles of governments in the economy
- More realistic, less deterministic theories of capitalism
- Amnesia about free time, the “forgotten American Dream” (Hunnicutt 2013)?

What Hasn't Changed as Much?

- Skidelsky and Skidelsky's three factors are still important:
  - Wants vs. needs
  - Pressure to work
  - Enjoyment from work
- The terms of the debate on jobs, technology, etc. (Bix 2001)
- Very few socio-technologically plausible scenarios are being seriously debated, let alone implemented

Long-term Goals in/for AI

- Russell/Norvig: rational action (vs. rational thinking, human-like thinking, or human-like action).
  - Autonomy as (a) goal.
- Human-centered computing, human-machine symbiosis, positive computing, etc. (Licklider 1950; Hoffman et al. 2012; Calvo and Peters 2014)
- Are these complementary? Mutually exclusive? Too soon to say?

Long-term Goals in/for AI (cont'd)

- Different social (Hoffman et al. 2012) and economic (Brynjolfsson and McAfee 2014) implications of different AI research decisions?
- Government funding priorities?

Design in/for AI

- User-friendliness
- Data, energy, etc. requirements
- Source code, demos, etc.
  - Replication, reuse, integration, collaboration, etc.
  - Democratization of AI

Design in/for AI (cont'd)

- Patent/licensing strategies viz-a-viz responsible innovation (Cooper 2013)
- Positive computing (Calvo and Peters 2014)

Applications

- Grand challenges
  - Health care
  - Sustainability
  - Etc.
- Are we thinking broadly/systematically enough?

Applications (cont'd)

- AI and inequality
  - Empowerment of elites, empowerment of everyone?
- What determines one outcome versus another?
  - Corporate involvement in AI

Work

- On the (dis)utility of work
  - Budd 2011: work is variously seen as curse, freedom, a commodity, occupational citizenship, disutility, personal fulfillment, a social relation, caring for others, identity, and service.
- End of work vs. end of bad work?

Work (cont'd)

- Quality and quantity of work in relation to AI goals/design decisions
- Technology not the key factor determining work quality (Morin 2004).
- Technical changes are insufficient to produce high quality work for all. Can help in some ways.

Education

- Convergent and divergent policy conclusions with respect to AI progress scenarios
- Convergent policies
  - High quality, accessible education for all
- Divergent policies
  - Vocational training?
  - Welfare?
  - Plausible career aspirations?
- Costs of different prediction errors

Leisure

- Leisure as common feature of utopias (Sargent 2010)
- Scarcity and cognitive bandwidth (Mullainathan/Shafir 2014)
  - Can AI help?
  - Disproportionate benefit hypothesis

Leisure (cont'd)

- AI for entertainment
- Post-(bad) work – leisure/work merge for more?
- Keynes: leisure society by ~2030.
    - Compare to AI progress theories, roadmaps, etc.
    - Not inevitable, but worth fighting for?

Basic Income

- AI progress and social progress
- Cost and quality of goods/services
- Basic income and work flexibility

Basic Income (cont'd)

- One AI-dependent basic income scenario:

  Piketty (capital tax) +

  Albus (ambitious, significant AI R+D funding) +

  Van Parijs (maximum sustainable basic income) +

  AI progress +

  Social progress and good governance =

- Progress toward post-(bad) work world by 2030?

AI Progress

- Environmental complexity (e.g. Russell/Norvig's framework) needs to be taken more seriously in the context of AI progress evaluation/anticipation
- No fact of the matter about which jobs will/won't be automated or how
- Technical AI progress versus diffusion of AI – complementary but not the same

AI Progress and Ethics

- Some AI progress/AI ethics connections:
  - Timeframes
  - Politics of expertise, hype, etc.
  - Public understanding of science
  - Policy decision-making (e.g. education)

AI Progress and Ethics (cont'd)

- Beyond Turing Test workshop
- Synthesis and analysis in/of AI
  - Progress theories/models as a bridge between the two (Kim 1990)
- Robust anticipation/shaping of change (Lempert et al. 2003)

Engaging the Public about AI

- Some models
  - Consensus conferences
  - Focus groups
  - Science fiction

Topics

- AI progress rates, drivers, future scenarios viz-a-viz policies for work, education, leisure, welfare, etc.
- Priorities/grand challenges
- Unexpected fears/hopes

Conclusions

- The world is different in some ways from when Keynes wrote his essay 1930, but not all.
- If we want to move to a post(-bad) work and/or leisure society, AI and robotics may help but are insufficient on their own.

Conclusions (cont'd)

- AI progress needs to be more robustly theorized and anticipated if we are to make good policy decisions.
- The public needs to know more and have more of a say in all these issues.

Acknowledgments

- Thanks to an anonymous reviewer, David Guston, Joanna Bryson, Erik Fisher, Daniel Dewey, Katja Grace, Tony Barrett, Deborah Strumsky, Beau Cronin, and Adam Elkus for helpful feedback on earlier versions of these ideas.
- This work was supported by the National Science Foundation under award #1257246 through the Virtual Institute of Responsible Innovation (VIRI). The findings and observations contained in this presentation are those of the author and do not necessarily reflect the views of the National Science Foundation.

Thanks!

www.milesbrundage.com

Miles.brundage@asu.edu