

A Research Agenda for Assessing the Economic Impacts of Code Generation Models

Sam Manning^{1‡}, Pamela Mishkin^{*2‡}, Gillian Hadfield³, Tyna Eloundou², and Emily Eisner⁴

¹*OpenResearch*

²*OpenAI*

³*University of Toronto*

⁴*University of California, Berkeley*

[‡]*These authors contributed equally to this work.*

March 3, 2022

Executive Summary

OpenAI is developing a research program to assess the economic impacts of code generation models and is inviting collaboration with external researchers. Rapid advances in the capabilities of large language models (LLMs) trained on code have made it increasingly important to study their economic impacts on individuals, firms, and society. Codex – an LLM developed by OpenAI by fine-tuning GPT-3 on billions of lines of publicly available code from GitHub – has been shown to generate functionally correct code 28.8% of the time on a sample of evaluation problems (Chen et al. 2021). This may have important implications for the future of coding and the economics of the industries that depend on it. In this document, we lay out a research agenda to assess the effects of Codex on economic factors of interest to policymakers, firms, and the public. We make a case for this research agenda by highlighting the potentially broad applicability of code generation models to software development, the potential for other LLMs to create significant social and economic impact as model capabilities advance, and the value of using Codex to generate evidence and establish methodologies that may be applicable to research on the economic impacts of future models. We propose that academic and policy research focus on studying code generation models and other LLMs so that evidence on their economic impacts can be used to inform decision-making in three key areas: Deployment policy, AI system design, and public policy. To help guide this research, we outline six priority outcome areas within the realm of economic impacts that we intend to use Codex to study: Productivity, Employment, Skill Development, Inter-firm

*Corresponding author, econ@openai.com.

Competition, Consumer Prices, and Economic Inequality. For each area, we briefly discuss previous literature on the impacts of artificial intelligence on each of these outcomes, describe questions that we believe to be key inputs to the three decision-making areas mentioned above, and provide examples of research that could be conducted with Codex. To catalyze work that builds off of this initial research agenda, we are announcing a [Call for Expressions of Interest](#) from external researchers to collaborate with OpenAI researchers and customers to better measure the economic impacts of code generation models and other LLMs.

1 Introduction

OpenAI is building out a research program to assess the economic impacts of code generation models with the goal of developing tools, methods, and partnerships that can enable improved research on the economic impacts of powerful language models. As code generation models and other large language models (LLMs) improve, they have the potential to impact many aspects of society, including work, productivity, skill development, and other economic outcomes. The depth and scope of the effects of code-generating LLMs will depend on how widespread their use becomes, which in turn depends on factors such as their capabilities and limitations, ease of use, associated costs, and the regulatory and institutional environments in which they are deployed. The capabilities of present and future code generation models may complement and/or substitute for the tasks completed by workers in coding-centric occupations (engineers, data analysts, software developers, etc.) by, for example:

- Impacting the costs associated with coding tasks
- Impacting the relative productivity of capital versus labor in the production process
- Shifting the allocation of tasks in the production process to capital vs labor
- Impacting the demand for existing skills (coding-centric and not) and spurring demand for new skills

These potential impacts are complex. Therefore, the research community’s ability to generate decision-relevant evidence on any of the research questions outlined in this document will be greatly enhanced by developing a range of productive partnerships, and we firmly believe that AI developers need to support external researchers undertaking this work, rather than conduct this research exclusively in-house. We hope this document serves as a starting point for collecting input from researchers, AI developers, policymakers, workers, labor unions, and firms interested in understanding the impacts of code generation models – and LLMs broadly – on economic outcomes. In Section 4 and in Table 1 below we highlight six research focus areas and key questions where OpenAI is interested in better understanding the economic impacts of code

generation models via Codex - an LLM developed by OpenAI that translates natural language to code (Chen et al. 2021).¹ Finally, we are issuing a [Call for Expressions of Interest](#) for external researchers to collaborate with OpenAI to better measure the economic impacts of code generation models, with the goal of building research methods and infrastructure that can be applied to other LLMs in the future. Similarly, we invite others deploying or using LLMs for code generation to support this work.

1.1 Call for Expressions of Interest

We are seeking feedback on this research agenda, as well as expressions of interest from individuals who are interested in partnering with OpenAI to study the economic impacts of Codex and to advise future research efforts on the economic impacts of novel LLMs. We welcome research proposals from all social science disciplines, including but not limited to economics, labor studies, sociology, and political science. We are also interested in engagement with private companies who have already integrated Codex. If you or your organization have a proposal for a research collaboration or would be interested in helping guide how OpenAI thinks about these issues, please see the link above for details on how to submit an expression of interest.

2 Motivations

2.1 Consider economic impacts as part of the AI Safety framework

A key motivation for the research agenda we propose in this paper is to ensure AI safety: even though the current capabilities of Codex do not threaten large-scale economic disruption or harm to human systems, future capabilities of code generation or other LLMs could. It is critical to engage in research about the economic impacts of model capabilities today in order to be positioned to assess the safety of developing and releasing more advanced systems in the future.

Foundational work setting the technical AI safety research agenda by Amodei, Olah, and coauthors has focused on the problem of "accidents in machine learning systems," while strongly supporting further work on privacy, security, fairness, economics, and policy (Amodei et al. 2016). The authors highlight the policy question "How do we predict and respond to the economic and social consequences of ML?" recognizing it as an important area, overlapping with other technical AI safety concerns, that warrants dedicated research. While far from the only such example, socioeconomic impacts are increasingly relevant as AI systems see increased adoption in and interaction with society (Weidinger et al. 2021).

¹This document does not present a comprehensive list of all potential areas of economic impact that would benefit from further research. This research agenda is an initial attempt at priority-setting given the range of critical questions on the economic impact of code generation models, and we are eagerly seeking feedback on what those priorities should be.

Direct Impacts & Priority Subquestions		
Research Area	Subquestions	Examples
Productivity	<ul style="list-style-type: none"> • What is the impact of Codex adoption on firm, team, and worker productivity? • What are the firm, worker, and use-case characteristics that drive differential impacts on productivity? • What are the mechanisms through which productivity impacts on firms, teams, and workers are realized? 	<ul style="list-style-type: none"> • Random assignment of model across workers, teams, and/or firms to assess impact on productivity-related outcomes • Longitudinal study of the production process as Codex applications are adopted and developed over time • Cataloging of products and projects built using Codex
Employment	<ul style="list-style-type: none"> • What is the impact of Codex adoption on the demand for human coding labor? • What is the impact of Codex adoption on the demand for human labor in non-coding roles? • What human coding tasks are most likely to be substituted by Codex and how is that labor reallocated? • What new tasks does Codex introduce into the production process and what skills are demanded to complete them? • What is the impact of Codex adoption on job quality? 	<ul style="list-style-type: none"> • Development of better benchmark datasets that map job tasks to model capabilities • Random assignment of model across workers, teams, and/or firms to assess impact on labor demand and job quality • Longitudinal study of team structure and labor demand as Codex applications are adopted and developed over time • Monitoring of job postings for tasks requiring proficiency with Codex or complementary skills
Skill Development	<ul style="list-style-type: none"> • How does the introduction of Codex to coding education programs change the skills that learners develop? • How does the adoption of Codex for use by advanced coders impact their coding innovation, creativity, and skill development? • What non-coding skill development trends are affected most by the applications built using the Codex API? • What implications does the use of Codex in education and training have for amplification of certain coding practices? 	<ul style="list-style-type: none"> • Qualitative data collection on the impact of Codex introduction to coding education programs on learning outcomes • Random assignment of model across workers, teams, and/or firms to assess impact on coding and non-coding skill development
Indirect Impacts & Priority Subquestions		
Research Area	Subquestions	Examples
Consumer Prices	<ul style="list-style-type: none"> • What is the impact of Codex adoption on the price of goods and services produced by the adopting entity? • What mechanisms drive observed impacts on prices, and how might these impacts scale with model improvements? 	<ul style="list-style-type: none"> • Development of an empirical framework for assessing the impact of code generation models on consumer prices
Inter-firm Competition	<ul style="list-style-type: none"> • What is the impact of Codex adoption on firm growth? How is this impact mediated by firm, industry and use-case characteristics? • Under what circumstances might Codex adoption increase the risk of harmful monopolies? 	<ul style="list-style-type: none"> • Identification of the firm and use-case characteristics that are likely to correlate with accelerated growth due to Codex adoption • Development of an empirical framework for assessing the impact of code generation models on intra-firm competition
Economic Inequality	<ul style="list-style-type: none"> • How does Codex adoption correlate with indicators of economic opportunity at the firm level (industry type, firm size, location, etc.) and individual level (income, wealth, race, gender, skills, zip code, etc.) • How can alternate deployment strategies reduce the risk of harmfully exacerbating economic inequalities? • How does Codex adoption change labor demand across the income and skill distribution? 	<ul style="list-style-type: none"> • Analysis of firm characteristics for firms that do and don't adopt Codex • Development of an empirical framework for assessing the impact of code generation models on income and wealth distributions • Monitoring and analyzing the evolution of wages across firms that do and don't adopt Codex (random assignment possible)

Table 1: Research focuses, key questions, and examples of research to collect evidence on economic impacts.

Systematic explorations of what might be considered “socio-economic safety” of models—the potential impacts of powerful AI systems on people and society as they interact with existing economic, social, and political institutions— may yield insights that are valuable to policymakers.

Absent policy intervention, LLMs may result in socio-economic safety risks by causing sudden negative impacts on the demand for human labor, increasing the frequency of labor market transitions, and exacerbating inequality, for example. Job displacement is associated with a range of negative impacts, including subsequent unemployment, long-term earnings losses, reduced psychological and physical well-being, family disruption, and lower levels of children’s educational attainment and well-being (Brand 2015, Young 2012, Schmillen 2020). Beyond affecting individual outcomes, economic impacts have the potential to shape the societal risk landscape in important ways. For example, at a societal level, sharp changes in the demand for human labor have been linked to higher levels of social unrest (Caprettini and Voth 2020). Depending on the fungibility of skills for those who experience a reduction in labor market opportunities as a result of AI system deployment, increasingly capable models risk exacerbating wage inequality, which in turn can amplify societal cleavages (Acemoglu and Restrepo 2021, Van de Werfhorst and Salverda 2012). In addition, differential access to required inputs to powerful LLMs – such as hardware, internet access, and digital literacy – will also perpetuate economic inequities (Weidinger et al. 2021). We must take these risks seriously and consider the potential implications for socio-economic safety when crafting deployment strategies and complimentary public policy proposals aimed at promoting well-being.

2.2 Incorporate economic impacts as inputs to key decisions

A central motivation for measuring economic impacts is to help researchers, firms, policymakers and the public better understand the populations most likely to benefit and those that could be negatively impacted from the adoption of AI systems that leverage LLMs. By better understanding the ways in which code generation models like Codex can impact economic outcomes for various actors in society, we can help inform decision-making in the three areas listed below.

- *Deployment policy*: Projected economic impacts are one of many criteria AI developers can use to inform if, when, and how a new system should be deployed to users and potential beneficiaries. By developing a deeper empirical understanding of the economic impacts of code generation models, research in this area can drive improved deployment policy that considers economic well-being as a key outcome.
- *AI system design*: Building our collective understanding of how a model like Codex can have tangible impacts on outcomes like productivity, employment, and skill development can illuminate ways in which future models can be designed for greater positive economic impact and fewer harms.

- *Public policy*: Research on the outcomes described in this agenda can identify potential economic impacts for which public policy intervention may be a helpful tool to improve economic outcomes and mitigate inequities that could be the product of the deployment of increasingly capable AI systems. A core goal of this stream of research is to generate improved data and produce novel evidence that can inform the policymaking process.

2.3 Build a test case for future research on the economic impacts of language models

The research that will be immediately shaped by this agenda will focus on the economic impacts of Codex, but we expect this research agenda to serve as a starting point for economic impacts research that can be applied more generally for future AI systems. There have been rapid advances in language model capabilities over the past several years (Brown et al. 2020, Dhariwal et al. 2020, Rae et al. 2022, Smith et al. 2022, Radford et al. 2021, Sun et al. 2021) and we recognize that as this progress continues, there will be a heightened need to carefully understand the evolution of economic impacts and translate this research into forecasting capabilities for new models. By articulating and executing on this research agenda via Codex, we aim to identify gaps in our approach, build research partnerships, solicit feedback, collect data on economic outcomes, and establish learning priorities that improve our collective ability to conduct policy-relevant economic impacts research on increasingly powerful language models in the future. The success of this agenda rests on the collaboration of the AI research community, policymakers, economists, and workers and we welcome your input.

2.4 Ensure that the economic impacts of progress towards AGI are broadly beneficial to humanity

OpenAI’s mission is to ensure that artificial general intelligence (AGI) – defined in OpenAI’s charter as “highly autonomous systems that outperform humans at most economically valuable work” – benefits all of humanity (OpenAI 2018). An important tenet of OpenAI’s deployment philosophy and policy is understanding and mitigating the safety risks of powerful AI models before deployment. If successful, highly capable autonomous systems are not only expected to transform the nature and quality of many jobs, but also perhaps engender structural economic changes, with impacts on inequality and employment. Previous major technological shifts such as the industrial revolution had positive long-run effects on many facets of economic life, yet they also caused economic hardship for segments of society that were affected by negative labor market shocks (Frey 2019). Therefore, it is critical that we generate evidence on the nature and distribution of impacts of new AI systems to ensure that their development and deployment can promote broad benefit to humanity in the short, medium, and long term.

3 What is Codex?

The economic impacts we will focus on in this research agenda are relevant to code generation models broadly. However, we plan to leverage OpenAI's Codex model to execute on this research agenda in the near-term. Codex is an example of an LLM - an artificial intelligence model trained to predict text to follow a given string of input text. For example, if an LLM like OpenAI's GPT-3 is given the prompt "I like to eat pizza because", it might generate the text "it is delicious." Codex is a fine-tuned version of OpenAI's GPT-3, meaning that it inherits GPT-3's language capacity and is given additional training on a wide range of programming languages (Brown et al. 2020, Chen et al. 2021). Its capabilities in natural language give it a remarkable ability to generalize to a wide range of tasks associated with coding, including code generation, code completion, code repair, code translation and code question answering. These capabilities have made it useful for a range of practical tasks, including generating code from natural language descriptions, writing documentation or unit tests for code snippets, completing partially written code, writing explanations for code snippets and fixing bugs in code. The model also has important limitations, namely that it often produces insecure code, can produce code that is not aligned with what the user intended, and is susceptible to reproducing or amplifying biases in the training data (Chen et al. 2021).

```
def fibonacci(n):
```

a.

```
def fibonacci(n):  
    """ This
```

b.

```
def fibonacci(n):  
    if n == 0:  
        return 0  
    elif n == 1:  
        return 1  
    else:  
        return fibonacci(n-1) + fibonacci(n-2)
```

c.

```
def fibonacci(n):  
    """ This returns the nth value of the Fibonacci series """  
    if n == 0: return 0  
    elif n == 1: return 1  
    else: return fibonacci(n-1)+fibonacci(n-2)
```

d.

One may want to implement a function in code that finds the n th number in the Fibonacci sequence. To write such a function, one might start with a **prompt**: some text that Codex uses as input for its generation. *a* and *b* above are prompts that we passed to Codex, containing the function name and expected arguments. Codex took a turn and completed *a* into the snippet in *c* and completed *b* into the snippet in *d*.

Codex can be accessed via an API, which users can access directly or via other products built using the API. A prominent example of a Codex-based application is Github Copilot – a tool developed by GitHub and OpenAI to autocomplete code and generate code based on natural language comments. In addition to Codex's built-in capabilities, Copilot is ever-present in compatible programming environments, suggesting code completions throughout a session, and it has the ability to propose up to 10 suggested code completions if requested.

As Codex's capabilities evolve, and as more developers build on top of the API, it is likely that the available applications will also evolve. While these

applications will be designed and released by external parties, OpenAI will likely exert some control over the capabilities of the underlying Codex model. Therefore, the economic impacts of Codex depend on the model’s inherent capabilities, and how widely used its downstream applications become. Understanding the core aspects of Codex adoption is essential to identifying the mechanisms through which Codex could have observable economic impacts, particularly as OpenAI controls the levers of who is given access and for what use cases. Furthermore, studying the mechanisms of potential economic impacts is critical to ensuring that research at OpenAI and in the broader community prioritizes the most pressing questions, identifies blindspots where potential economic harms might exist, and makes evidence-based assumptions about how economic impacts may change as model capabilities evolve.

4 Research Agenda: Focus Areas

This section outlines several preliminary focus areas for our research agenda on the economic impacts of code generation models. We divide these focus areas into two categories:

1. Direct impacts, which will include productivity, employment, and skill development, and
2. Indirect impacts which will include inter-firm competition, consumer prices, and economic inequality.

The distinction between direct and indirect impacts is not meant to understate the importance of the indirect impacts as drivers of economic well-being. The categorization is useful to highlight the fact that research on direct impacts will often be a necessary input for precise research on indirect impacts. For example, to assess the impacts of code generation models on economic inequality, it is critical to better understand the distribution of impacts on employment and wages. Similarly, in order to enhance our understanding of how these models impact consumer prices, it is helpful to measure whether or not they introduce any changes in productivity within the production process for goods and services.

While this section identifies potential economic impacts of code generation models beyond just Codex, we plan to use Codex to generate evidence on the magnitude and direction of impacts. As such, we speak below about the potential impacts that Codex specifically may have on individuals, firms, and society.

The impacts of LLMs such as Codex on economic outcomes will vary widely depending on a number of underlying factors (Frank et al. 2019, Klinova and Korinek 2021, Trammell and Korinek 2021, Weidinger et al. 2021). Understanding the differential impact of code generation models – whether mediated by use-case, geography, labor market, firm, or individual characteristics – will be a priority for research across all of the focus areas described below.

4.1 Direct Impacts

4.1.1 Productivity

Background Neoclassical economic theory predicts that at the aggregate level, technological progress increases overall productivity (Romer 1990, Solow 1956). However, recent decades have not seen as strong productivity growth as might have been expected given rapid advancement in technology (Gordon 2018, Brynjolfsson, Rock, and Syverson 2017). In order to project the oncoming productivity impacts of AI, Brynjolfsson, Benzell, and Rock warn against relying on previous trends and instead suggest a need to “... study and understand the specific technologies that actually exist and make an assessment of their potential.” (Brynjolfsson, Benzell, and Rock 2020) The roll-out of Codex presents an opportunity to study the micro-level impact of code-generating AI on individual level productivity, a subject that will be key to understanding the current relationship between technological progress and economic growth.

Damioli and coauthors take a step in this direction by examining data from 5,257 firms worldwide that filed one or more patents related to AI between 2000 and 2016 (Damioli, Van Roy, and Vertesy 2021). The authors find that AI patent applications have a positive effect on within-firm labor productivity. This study is among the first to estimate a causal relationship between new AI technologies and the productivity of the firms that develop those technologies. Indeed, literature on the causal impact of AI on individual firms is scarce, largely due to a lack of firm-level data. Multiple recent papers make an explicit call for more firm-level data in order to build a clearer understanding of the impact of AI on a range of economic outcomes, and how those impacts are mediated by firm characteristics (Seamans and Raj 2018, Frank et al. 2019). Through OpenAI’s partnerships with firms that have adopted Codex, we intend to build on previous research that has used novel data collection approaches to measure the impact of code generation tools on productivity (Xu, Vasilescu, and Neubig 2021) and respond directly to this call for further firm-level data by examining the impact of Codex on both worker and firm-level measures of productivity.

How Codex May Impact Productivity Codex has the potential to increase the productivity of individual workers in coding-centric roles. The adoption of Codex could reduce the amount of time needed to look up syntax, reference old code, add documentation, write basic programs or switch between tasks and projects. Individuals who use Codex models or applications could also realize productivity effects via faster code, higher code quality, or improved documentation. Through the applications built with Codex, productivity could be enhanced not solely for coding tasks but for many tasks related to design, engineering, and data visualization. We are interested in understanding the distribution of productivity impacts on workers across the spectrum of tasks, skills and roles. This includes workers in coding-centric roles as well as workers in non-coding positions who may be affected by increased automation or adoption of productivity-enhancing tools built using Codex.

Broad Research Questions

- What is the impact of Codex adoption on firm, team, and worker productivity?
- What are the firm, worker, and use-case characteristics that drive differential impacts on productivity?
- What are the mechanisms through which productivity impacts firms, teams, and individual workers?

4.1.2 Employment

Background A growing literature in economics has renewed the recent focus of researchers on the potential impacts of technological advancement on employment (Acemoglu and Restrepo 2018, Autor 2015, Brynjolfsson and McAfee 2014, Mokyr, Vickers, and Ziebarth 2015, Tolan et al. 2021). Frey and Osborne estimate that 47% of total US employment is susceptible to automation (Frey and Osborne 2017). Aghion and coauthors highlight that the aggregate effects of AI on employment will be heavily mediated by competition, labor, and education policy (Aghion, Antonin, and Bunel 2020). Expert forecasts vary in their predictions, but overall suggest a considerable chance that AI will surpass human capabilities at most tasks within several decades.²

How Codex May Impact Employment The adoption of Codex and other code-generating AI could have a potentially large impact on employment in the technology and information sectors. As Codex’s capabilities continue to expand, Codex may eventually serve as a substitute for a larger share of coding tasks currently completed by human labor. Alternatively, Codex may augment human labor such that it is adopted as a net complement to labor and increases the demand for workers who perform tasks such as detailed code review, intensive quality assurance, or the application of sales and logistics expertise. Additionally, Codex could spark a need for new skills, changing team composition and shifting demand towards new tasks in which labor has a comparative advantage, a phenomenon researchers have called the “reinstatement effect” (Acemoglu and Restrepo 2019). The effects of code generation models on the completion of micro-work tasks outsourced by firms to gig-economy workers is another potential avenue of impact on worker opportunity and well-being.

With respect to Codex, we are interested in empirically assessing how these dynamics will unfold, particularly as the model progresses in its capabilities. Understanding the balance of displacement versus reinstatement of tasks and jobs across different industries, firms, and use-cases is an essential input to

²Expert forecasts collected by Grace and coauthors, for example, give a 50% chance that AI systems will outperform humans at all tasks by 2063, and a 10% probability that those capabilities will exist by 2027 (Grace et al. 2018). More recent forecasts collected by Gruetzemacher and coauthors suggest there is a 50% chance that AI systems would be capable of automating 90% of human tasks by 2045.(Gruetzemacher, Paradise, and Lee 2020)

forecasting future direct labor market impacts as the capabilities of Codex and other code-generating models evolve.

Of particular interest is whether we can leverage worker and firm-level data to identify trends in the potential demand shifts for various types of skills and how fungible those skills are in the labor market. If we expect Codex to drive down demand for entry level coders (or other roles with rote and repetitive coding tasks) but drive up demand for senior engineers and managers, for example, then we will want to have an informed estimate of the impacts that may have on wage and mobility outcomes to inform deployment and public policy decisions. We hope that foundational research on the employment impacts of Codex can enable increasingly policy-relevant research to be done to project longer-term impacts of future code-generating AI models.

In addition to impacts on total employment, Codex may also impact job quality and the nature of work itself. Broadly, advances in AI have the potential to reduce occupational safety risks for certain jobs, create new opportunities for aging workers or those with disabilities, and substitute for overly repetitive and mundane tasks (EU-OSHA 2021). However, increased automation can drive social isolation at work, increased specialization, performance pressure, reduced worker autonomy and overbearing worker surveillance, all of which may reduce well-being on the job (Kaplan and Schulhofer-Wohl 2018, Partnership on AI 2020, Weidinger et al. 2021). Measuring the effects of Codex on job quality is a key input to understanding the broader impacts of Codex on worker well-being.

Potential Research Questions

- What is the impact of Codex adoption on the demand for human coding labor?
- What is the impact of Codex adoption on the demand for human labor in non-coding roles?
- What human coding tasks are most likely to be substituted by Codex and how is that labor reallocated?
- What new tasks does Codex introduce into the production process and what skills are demanded to complete them?
- What is the impact of Codex adoption on job quality?

4.1.3 Skill development

Background A large body of literature suggests that complementarities between technological advances and high-skilled labor can drive increasing returns to skill development (Acemoglu and Autor 2011, Bound and Johnson 1992, Goos 2018, Katz and Murphy 1992). Predictable pathways towards a labor reinstatement effect from Codex include increased demand for skills such as prompt engineering, Codex-specific debugging, and specialized quality assurance of AI-generated outputs. Given the likelihood that Codex could generate demand

for new skills in the labor force, we would like to examine the ways that Codex can also drive the development of new skills when incorporated into training and education programs. By examining this question empirically with Codex, we intend to contribute to a body of literature that has investigated the impact of technological development on skill development. Several descriptive case studies summarize the experiences of students or firms that integrate low-code software tools into work and learning environments (Beranic, Rek, and Hericko 2020, Corral, Fronza, and Pahl 2021). However, we are not aware of any empirical work estimating the impact of these tools on skill development or retention.

How Codex May Impact Skill Development The ability for Codex to make coding suggestions could either enhance a user’s learning process or create inattentive reliance on Codex that may stifle creativity and iterative learning. It is plausible that Codex suggestions disincentivize coders from learning or retaining new knowledge when they feel they can rely on Codex. We are particularly interested in learning whether or not this is the case at the frontier of human coding innovation and skill development. Estimating the impact of Codex on coding skill development can help us understand the impact on human coding innovation – an important driver of technological progress and an essential data input for increasingly powerful code generation tools. Furthermore, evaluating the impacts of Codex on skill development for coders and non-coders alike can influence decisions about future education policy and the design of training programs that match the needs of the economy.

Potential Research Questions

- How does the introduction of Codex to coding education programs change the skills that learners develop?
- How does the adoption of Codex for use by advanced coders impact their coding innovation, creativity, and skill development?
- What non-coding skill development trends are impacted most by the applications built using the Codex API?
- What implications might the use of Codex in education and training have for amplification of certain coding practices?

4.2 Indirect Impacts

The outcomes included in this section are listed separately from those above purely because we expect the outputs from research on the “direct” impacts above to be key inputs into understanding the impact of Codex on these “indirect” impacts. The distinction between direct and indirect impacts does not reflect a difference in the relative importance of the outcomes in either group within this research agenda.

4.2.1 Consumer Prices

Background Technological progress has made the production of countless goods and services cheaper over time (Roser 2016). Researchers have speculated that as the general capabilities of AI advance, the costs of labor to produce many goods and services could fall dramatically, driving a reduction in the market price for consumer goods and services (Stone et al. 2016). Such an impact would rely on AI systems introducing productivity and efficiency gains into the production process, including by substituting human labor with automated systems that run at lower marginal costs.

How Codex May Affect Consumer Prices Codex provides a tangible opportunity to better understand how the introduction of a specific, potentially powerful AI system can impact the costs of production, and how that impact is passed on to consumers via prices. By augmenting any production process that in part relies on code generation, Codex could have a downstream impact on the prices of goods and services. Through partnerships with firms that have adopted Codex, we can learn about the impact of Codex on factors of production, and begin to build an understanding of how those impacts are passed on to consumers, if at all. Given the growing importance of coding and software as an input to the production of goods and services, understanding this impact for one code generation model could foster better understanding of the potential impacts of increasingly capable code generation models in the future.

Potential Research Questions

- What is the impact of Codex adoption on the price of goods and services produced by the adopting entity?
- What mechanisms drive observed impacts on prices, and how might these impacts scale with model improvements?

4.2.2 Inter-firm competition

Background AI-adopting firms with a better ability to collect and use data – specifically data that is inaccessible to their competitors – may drive “unfair competition” (Acemoglu 2021a). As a result, particularly well-positioned firms could capture excessive consumer surplus and relax price competition in the market (Acemoglu 2021a). Investments in AI technology have been shown to be correlated with increased firm growth, particularly among already large firms relative to others in their industry (Babina et al. 2021). Better understanding the potential for Codex to drive increased industry concentration is a critical input to improved deployment strategy and public policy design.

How Codex May Impact Inter-firm Competition The effective adoption of Codex could spark productivity and efficiency gains, potentially driving faster growth at the firm level. We’re interested in understanding the characteristics of

a firm that make it more likely to realize the economic impacts from Codex. Are there existing monopolies within industries that Codex would further entrench? What impact would the adoption of Codex have on competition and what role should those impacts play in deployment policy?

A deeper understanding of the impacts of modern AI-system adoption on competition is urgently needed. However, without a sample of several hundred firms, many confounding factors would limit our ability to causally identify the impact of Codex on firm-level competition dynamics. As such, our priority in the short term is to enhance our understanding of the mechanisms through which Codex might accelerate firm-level growth, focusing empirical research on the “direct” impacts described previously in this document that might effect market dynamics. We encourage expressions of interest from scholars interested in guiding our approach to better understanding impacts on competition dynamics and how Codex might impact the underlying drivers of shifts in market power.

Potential Research Questions

- What is the impact of Codex adoption on firm growth? How is this impact mediated by firm and industry characteristics?
- Under what circumstances might Codex adoption increase the risk of harmful monopolies?

4.2.3 Economic Inequality

Background In the US, the average 2021 annual income among individuals in the top 1% of earners (\$1.6m) was approximately 84x higher than the average income of individuals in the bottom 50% of earners (\$19.1k) (Blanchet, Saez, and Zuckman 2022). The divergence of both income and wealth in the US since the 1980s has been attributed in part to the economic impacts of technological change (Jaumotte, Lall, and Papageorgiou 2013, Acemoglu 2002, Rotman 2014). Numerous studies have demonstrated that middle-wage jobs have been increasingly displaced through technological innovation in recent decades. Highly routine jobs have been particularly susceptible to displacement, while those requiring abstract or manual tasks (professional, managerial, and technical occupations at the higher end of the wage spectrum as well as service and labor jobs at the other) have proven less susceptible (Autor 2015, Autor, Levy, and Murnane 2003, Autor and Dorn 2013, Goos and Manning 2007). This phenomenon has been termed “job polarization” and has been attributed to skill-biased and routine-biased technological change (Berman, Bound, and Machin 1998, Goos and Manning 2007, Goos, Manning, and Salomons 2014). A core driver of the distributive economic impacts of LLMs and other AI systems is whether they are primarily used to augment and complement human labor or replace it (Brynjolfsson 2022, Acemoglu and Restrepo 2021).

How Codex May Affect Economic Inequality Codex presents an example of how the scope of “routine” automatable tasks can change over time (Lu 2015).

This shift may be gradual and uneven, particularly across different labor markets, with some workers and firms adopting new technologies more readily than others. This may lead to a widening of existing disparities in skill, training, or digital literacy, or to greater inequality in the distribution of economic benefits from technology.

The adoption of new technologies and automation methods is not inevitable. Different firms and workers may have different preferences and costs for adopting new technology. In addition, some workers may be unable to adopt new technologies due to the high cost of complementary technologies, the high cost of retraining, or insufficient digital literacy. The adoption of Codex therefore may correlate with – and exacerbate – existing inequities in technology access, digital literacy, and economic opportunity. There is a risk that the economic benefits of code generation models may be shared unequally, with much of the gains flowing to the owners of capital, such as investors and shareholders.

By partnering with external academics and Codex customers, we aim to foster research that helps assess the impact of Codex on the distribution of income, skills, wealth, and economic mobility. The outcomes of this research will be key inputs to policy design aimed at mitigating any distributional impacts of new AI systems that may amplify harmful inequities.

Potential Research Questions

- How does Codex adoption correlate with other indicators of economic opportunity and mobility at the firm level (industry type, firm size, location, etc.) and individual level (income, wealth, race, zip code, etc.)?
- How can alternate model deployment strategies reduce the risk of harmfully exacerbating economic inequalities?
- How does Codex adoption change labor demand across the income and skill distribution?

5 Prioritization

We listed numerous avenues for research above and we encourage collaborations to pursue them all. When considering which projects to initiate, we will prioritize research that has the following characteristics:

- Helps build sustained partnerships for data sharing and research collaboration that can improve learning about the economic impacts of LLMs over time.
- Has the potential to inform deployment decisions for code generation models or could directly influence public policy decisions meant to enhance the economic benefits of these models and minimize any negative impacts.

- Helps segment aspects of code generation models based on their likely economic impact, both positive and negative, in order to inform future model design decisions.
- Helps OpenAI, other AI developers and external research partners estimate the potential future economic impacts of improved code generation models.
- Is unlikely to happen without OpenAI support.
- Is most likely to succeed if led by researchers who are external to OpenAI.

6 Conclusion

This research agenda is just one of several recent contributions meant to inform the direction of future work to ensure that the economic impacts of AI are as universally positive as possible (Acemoglu 2021a, Acemoglu 2021b, Partnership on AI 2021, Siddarth et al. 2021, Weidinger et al. 2021, Autor, Mindell, and Reynolds 2022). We are excited by progress in the fields of AI ethics, safety, and alignment research and recognize that as the capabilities of AI systems advance, so too will the potential impacts of key decisions related to AI system design, deployment, and public policy. It is our hope that this research agenda will not only inspire deeper conversation about the economic impacts of increasingly capable LLMs but also – paired with the [Call for Expressions of Interest](#) – catalyze concrete action to measure economic impacts and inform decision-making in these areas.

Call for Expressions of Interest If you are a researcher interested in partnering with OpenAI researchers and customers to study the economic impacts of Codex, please see the link above to read more and for details on how to submit an expression of interest.

Acknowledgements Thanks to Steven Adler, Lama Ahmad, Stephanie Bell, Miles Brundage, Katya Klinova, Gretchen Krueger, Jade Leung, Anna Mankanju, Katie Mayer, Richard Ngo, Cullen O’Keefe, Girish Sastry, Sarah Shoker, and Natalie Staudacher for feedback on drafts of this document. Thanks to Michelle Alexopoulos, Sarah Bana, Alex Bartik, Erik Brynjolfsson, Tim de Stefano, Avi Goldfarb, Marlène Koffi, Mina Lee, Zanele Munyikwa, Mark Muro, Frank Nagle, Maria del Rio-Chanona, Daniel Rock, Anna Salomons, and Ben Weidmann for helpful discussions on potential avenues for research on the economic impacts of code generation models.

References

Acemoglu, Daron (2002). “Technical Change, Inequality, and the Labor Market”. In: *Journal of Economic Literature* 40.1, pp. 7–72. ISSN: 0022-0515.

- Acemoglu, Daron (Sept. 2021a). *Harms of AI*. Tech. rep. w29247. Cambridge, MA: National Bureau of Economic Research, w29247. DOI: [10.3386/w29247](https://doi.org/10.3386/w29247).
- ed. (2021b). *Redesigning AI: Work, Democracy, and Justice in the Age of Automation*. Boston Review/Forum 18 (46.2). Cambridge, MA: Boston Review. ISBN: 978-1-946511-62-1.
- Acemoglu, Daron and David Autor (2011). “Skills, Tasks and Technologies: Implications for Employment and Earnings”. In: *Handbook of Labor Economics*. Vol. 4. Elsevier, pp. 1043–1171. ISBN: 978-0-444-53452-1. DOI: [10.1016/S0169-7218\(11\)02410-5](https://doi.org/10.1016/S0169-7218(11)02410-5).
- Acemoglu, Daron and Pascual Restrepo (June 2018). “The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment”. In: *American Economic Review* 108.6, pp. 1488–1542. ISSN: 0002-8282. DOI: [10.1257/aer.20160696](https://doi.org/10.1257/aer.20160696).
- (May 2019). “Automation and New Tasks: How Technology Displaces and Reinstates Labor”. In: *Journal of Economic Perspectives* 33.2, pp. 3–30. ISSN: 0895-3309. DOI: [10.1257/jep.33.2.3](https://doi.org/10.1257/jep.33.2.3).
- (June 2021). *Tasks, Automation, and the Rise in US Wage Inequality*. Tech. rep. w28920. Cambridge, MA: National Bureau of Economic Research, w28920. DOI: [10.3386/w28920](https://doi.org/10.3386/w28920).
- Aghion, Philippe, Céline Antonin, and Simon Bunel (Jan. 2020). “Artificial Intelligence, Growth and Employment: The Role of Policy”. In: *Economie et Statistique / Economics and Statistics* 510-511-512, pp. 149–164. ISSN: 03361454. DOI: [10.24187/ecostat.2019.510t.1994](https://doi.org/10.24187/ecostat.2019.510t.1994).
- Amodei, Dario et al. (July 2016). “Concrete Problems in AI Safety”. In: *arXiv:1606.06565 [cs]*. arXiv: [1606.06565 \[cs\]](https://arxiv.org/abs/1606.06565).
- Autor, David (Aug. 2015). “Why Are There Still So Many Jobs? The History and Future of Workplace Automation”. In: *Journal of Economic Perspectives* 29.3, pp. 3–30. ISSN: 0895-3309. DOI: [10.1257/jep.29.3.3](https://doi.org/10.1257/jep.29.3.3).
- Autor, David and David Dorn (Aug. 2013). “The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market”. In: *American Economic Review* 103.5, pp. 1553–1597. ISSN: 0002-8282. DOI: [10.1257/aer.103.5.1553](https://doi.org/10.1257/aer.103.5.1553).
- Autor, David, Frank Levy, and Richard J Murnane (2003). “The Skill Content of Recent Technological Change: An Empirical Exploration”. In: *Quarterly Journal of Economics*.
- Autor, David, David A. Mindell, and Elisabeth B. Reynolds (2022). *The Work of the Future: Building Better Jobs in an Age of Intelligent Machines*. The MIT Press. ISBN: 978-0-262-36775-2. DOI: [10.7551/mitpress/14109.001.0001](https://doi.org/10.7551/mitpress/14109.001.0001).
- Babina, Tania et al. (Nov. 2021). “Artificial Intelligence, Firm Growth, and Product Innovation”.
- Beranec, Tina, Patrik Rek, and Marjan Hericko (Oct. 2020). “Adoption and Usability of Low-Code/No-Code Development Tools”. In: *Proceedings of the Central European Conference on Information and Intelligent Systems*. Varazdin, Croatia.
- Berman, Eli, John Bound, and Stephen Machin (Nov. 1998). “Implications of Skill-Biased Technological Change: International Evidence*”. In: *The Quarterly*

- Journal of Economics* 113.4, pp. 1245–1279. ISSN: 1531-4650, 0033-5533. DOI: [10.1162/003355398555892](https://doi.org/10.1162/003355398555892).
- Blanchet, Thomas, Emmanuel Saez, and Gabriel Zuckman (Feb. 2022). *Realtime Inequality*. <https://realtimeinequality.org/>.
- Bound, John and George Johnson (1992). “Changes in the Structure of Wages in the 1980’s: An Evaluation of Alternative Explanations”. In: *The American Economic Review* 82.3, pp. 371–392. ISSN: 00028282.
- Brand, Jennie E. (Aug. 2015). “The Far-Reaching Impact of Job Loss and Unemployment”. In: *Annual Review of Sociology* 41.1, pp. 359–375. ISSN: 0360-0572, 1545-2115. DOI: [10.1146/annurev-soc-071913-043237](https://doi.org/10.1146/annurev-soc-071913-043237).
- Brown, Tom B. et al. (July 2020). “Language Models Are Few-Shot Learners”. In: *arXiv:2005.14165 [cs]*. arXiv: [2005.14165 \[cs\]](https://arxiv.org/abs/2005.14165).
- Brynjolfsson, Erik (Jan. 2022). “The Turing Trap: The Promise & Peril of Human-Like Artificial Intelligence”. In: *arXiv:2201.04200 [cs, econ, q-fin]*. arXiv: [2201.04200 \[cs, econ, q-fin\]](https://arxiv.org/abs/2201.04200).
- Brynjolfsson, Erik, Seth Benzell, and Daniel Rock (2020). *Understanding and Addressing the Modern Productivity Paradox*. Research Brief. MIT.
- Brynjolfsson, Erik and Andrew McAfee (2014). *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. First edition. New York: W.W. Norton & Company. ISBN: 978-0-393-23935-5.
- Brynjolfsson, Erik, Daniel Rock, and Chad Syverson (Nov. 2017). *Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics*. Tech. rep. w24001. Cambridge, MA: National Bureau of Economic Research, w24001. DOI: [10.3386/w24001](https://doi.org/10.3386/w24001).
- Caprettini, Bruno and Hans-Joachim Voth (Sept. 2020). “Rage against the Machines: Labor-Saving Technology and Unrest in Industrializing England”. In: *American Economic Review: Insights* 2.3, pp. 305–320. ISSN: 2640-205X, 2640-2068. DOI: [10.1257/aeri.20190385](https://doi.org/10.1257/aeri.20190385).
- Chen, Mark et al. (July 2021). “Evaluating Large Language Models Trained on Code”. In: *arXiv:2107.03374 [cs]*. arXiv: [2107.03374 \[cs\]](https://arxiv.org/abs/2107.03374).
- Corral, Luis, Ilenia Fronza, and Claus Pahl (Oct. 2021). “Block-Based Programming Enabling Students to Gain and Transfer Knowledge with a No-code Approach”. In: *Proceedings of the 22st Annual Conference on Information Technology Education*. SnowBird UT USA: ACM. ISBN: 978-1-4503-8355-4. DOI: [10.1145/3450329.3478314](https://doi.org/10.1145/3450329.3478314).
- Damioli, Giacomo, Vincent Van Roy, and Daniel Vertesy (Mar. 2021). “The Impact of Artificial Intelligence on Labor Productivity”. In: *Eurasian Business Review* 11.1, pp. 1–25. ISSN: 1309-4297, 2147-4281. DOI: [10.1007/s40821-020-00172-8](https://doi.org/10.1007/s40821-020-00172-8).
- Dhariwal, Prafulla et al. (Apr. 2020). “Jukebox: A Generative Model for Music”. In: *arXiv:2005.00341 [cs, eess, stat]*. arXiv: [2005.00341 \[cs, eess, stat\]](https://arxiv.org/abs/2005.00341).
- Frank, Morgan R. et al. (Apr. 2019). “Toward Understanding the Impact of Artificial Intelligence on Labor”. In: *Proceedings of the National Academy of Sciences* 116.14, pp. 6531–6539. ISSN: 0027-8424, 1091-6490. DOI: [10.1073/pnas.1900949116](https://doi.org/10.1073/pnas.1900949116).

- Frey, Carl Benedikt (2019). *The Technology Trap: Capital, Labor, and Power in the Age of Automation*. First paperback printing. Princeton, New Jersey Oxford: Princeton University Press. ISBN: 978-0-691-21079-7 978-0-691-17279-8.
- Frey, Carl Benedikt and Michael A. Osborne (Jan. 2017). “The Future of Employment: How Susceptible Are Jobs to Computerisation?” In: *Technological Forecasting and Social Change* 114, pp. 254–280. ISSN: 00401625. DOI: [10.1016/j.techfore.2016.08.019](https://doi.org/10.1016/j.techfore.2016.08.019).
- Goos, Maarten (July 2018). “The Impact of Technological Progress on Labour Markets: Policy Challenges”. In: *Oxford Review of Economic Policy* 34.3, pp. 362–375. ISSN: 0266-903X, 1460-2121. DOI: [10.1093/oxrep/gry002](https://doi.org/10.1093/oxrep/gry002).
- Goos, Maarten and Alan Manning (Feb. 2007). “Lousy and Lovely Jobs: The Rising Polarization of Work in Britain”. In: *Review of Economics and Statistics* 89.1, pp. 118–133. ISSN: 0034-6535, 1530-9142. DOI: [10.1162/rest.89.1.118](https://doi.org/10.1162/rest.89.1.118).
- Goos, Maarten, Alan Manning, and Anna Salomons (Aug. 2014). “Explaining Job Polarization: Routine-Biased Technological Change and Offshoring”. In: *American Economic Review* 104.8, pp. 2509–2526. ISSN: 0002-8282. DOI: [10.1257/aer.104.8.2509](https://doi.org/10.1257/aer.104.8.2509).
- Gordon, Robert (Apr. 2018). *Why Has Economic Growth Slowed When Innovation Appears to Be Accelerating?* Tech. rep. w24554. Cambridge, MA: National Bureau of Economic Research, w24554. DOI: [10.3386/w24554](https://doi.org/10.3386/w24554).
- Grace, Katja et al. (July 2018). “Viewpoint: When Will AI Exceed Human Performance? Evidence from AI Experts”. In: *Journal of Artificial Intelligence Research* 62, pp. 729–754. ISSN: 1076-9757. DOI: [10.1613/jair.1.11222](https://doi.org/10.1613/jair.1.11222).
- Gruetzmacher, Ross, David Paradise, and Kang Bok Lee (Dec. 2020). “Forecasting Extreme Labor Displacement: A Survey of AI Practitioners”. In: *Technological Forecasting and Social Change* 161, p. 120323. ISSN: 00401625. DOI: [10.1016/j.techfore.2020.120323](https://doi.org/10.1016/j.techfore.2020.120323).
- Jaumotte, Florence, Subir Lall, and Chris Papageorgiou (June 2013). “Rising Income Inequality: Technology, or Trade and Financial Globalization?” In: *IMF Economic Review* 61.2, pp. 271–309. ISSN: 2041-4161, 2041-417X. DOI: [10.1057/imfer.2013.7](https://doi.org/10.1057/imfer.2013.7).
- Kaplan, Greg and Sam Schulhofer-Wohl (Aug. 2018). “The Changing (Dis-)Utility of Work”. In: *Journal of Economic Perspectives* 32.3, pp. 239–258. ISSN: 0895-3309. DOI: [10.1257/jep.32.3.239](https://doi.org/10.1257/jep.32.3.239).
- Katz, L. F. and K. M. Murphy (Feb. 1992). “Changes in Relative Wages, 1963-1987: Supply and Demand Factors”. In: *The Quarterly Journal of Economics* 107.1, pp. 35–78. ISSN: 0033-5533, 1531-4650. DOI: [10.2307/2118323](https://doi.org/10.2307/2118323).
- Klinova, Katya and Anton Korinek (July 2021). “AI and Shared Prosperity”. In: *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*. Virtual Event USA: ACM, pp. 645–651. ISBN: 978-1-4503-8473-5. DOI: [10.1145/3461702.3462619](https://doi.org/10.1145/3461702.3462619).
- Lu, Qian (2015). “The End of Polarization? Technological Change and Employment in the U.S. Labor Market”.
- Mokyr, Joel, Chris Vickers, and Nicolas L. Ziebarth (Aug. 2015). “The History of Technological Anxiety and the Future of Economic Growth: Is This Time

- Different?” In: *Journal of Economic Perspectives* 29.3, pp. 31–50. ISSN: 0895-3309. DOI: [10.1257/jep.29.3.31](https://doi.org/10.1257/jep.29.3.31).
- OpenAI (2018). *OpenAI Charter*. <https://openai.com/charter/>.
- EU-OSHA (2021). *Impact of Artificial Intelligence on Occupational Safety and Health: Policy Brief*.
- Partnership on AI (2020). *Framework for Promoting Workforce Well-being in the AI-Integrated Workplace*.
- (2021). *Redesigning AI for Shared Prosperity: An Agenda*.
- Radford, Alec et al. (Feb. 2021). “Learning Transferable Visual Models From Natural Language Supervision”. In: *arXiv:2103.00020 [cs]*. arXiv: [2103.00020 \[cs\]](https://arxiv.org/abs/2103.00020).
- Rae, Jack W. et al. (Jan. 2022). “Scaling Language Models: Methods, Analysis & Insights from Training Gopher”. In: *arXiv:2112.11446 [cs]*. arXiv: [2112.11446 \[cs\]](https://arxiv.org/abs/2112.11446).
- Romer, Paul M. (Oct. 1990). “Endogenous Technological Change”. In: *Journal of Political Economy* 98.5, Part 2, S71–S102. ISSN: 0022-3808, 1537-534X. DOI: [10.1086/261725](https://doi.org/10.1086/261725).
- Roser, Christoph (Oct. 2016). *Faster, Better, Cheaper in the History of Manufacturing: From the Stone Age to Lean Manufacturing and Beyond*. Zeroth. 1 Edition. | Boca Raton : CRC Press, 2016.: Productivity Press. ISBN: 978-1-315-36794-1. DOI: [10.1201/9781315367941](https://doi.org/10.1201/9781315367941).
- Rotman, David (Oct. 2014). “Technology and Inequality”. In: *MIT Technology Review*.
- Schmillen, Achim D. (May 2020). *Causes and Impacts of Job Displacements and Public Policy Responses*. Tech. rep. World Bank, Washington, DC. DOI: [10.1596/33720](https://doi.org/10.1596/33720).
- Seamans, Robert and Manav Raj (2018). “AI, Labor, Productivity and the Need for Firm-Level Data”. In: *National Bureau of Economic Research*.
- Siddarth, Divya et al. (2021). *How AI Fails Us*. Tech. rep. Justice, Health & Democracy Impact Initiative, Edmond J. Safra Center for Ethics, Harvard University.
- Smith, Shaden et al. (Feb. 2022). “Using DeepSpeed and Megatron to Train Megatron-Turing NLG 530B, A Large-Scale Generative Language Model”. In: *arXiv:2201.11990 [cs]*. arXiv: [2201.11990 \[cs\]](https://arxiv.org/abs/2201.11990).
- Solow, Robert M. (Feb. 1956). “A Contribution to the Theory of Economic Growth”. In: *The Quarterly Journal of Economics* 70.1, p. 65. ISSN: 00335533. DOI: [10.2307/1884513](https://doi.org/10.2307/1884513).
- Stone, Peter et al. (Sept. 2016). *Artificial Intelligence and Life in 2030. One Hundred Year Study on Artificial Intelligence: Report of the 2015-2016 Study Panel*. Tech. rep. Stanford, CA: Stanford University.
- Sun, Yu et al. (July 2021). “ERNIE 3.0: Large-scale Knowledge Enhanced Pre-training for Language Understanding and Generation”. In: *arXiv:2107.02137 [cs]*. arXiv: [2107.02137 \[cs\]](https://arxiv.org/abs/2107.02137).
- Tolan, Songül et al. (June 2021). “Measuring the Occupational Impact of AI: Tasks, Cognitive Abilities and AI Benchmarks”. In: *Journal of Artificial*

- Intelligence Research* 71, pp. 191–236. ISSN: 1076-9757. DOI: [10.1613/jair.1.12647](https://doi.org/10.1613/jair.1.12647).
- Trammell, Philip and Anton Korinek (2021). “Economic Growth under Transformative AI”.
- Van de Werfhorst, Herman G. and Wiemer Salverda (Dec. 2012). “Consequences of Economic Inequality: Introduction to a Special Issue”. In: *Research in Social Stratification and Mobility* 30.4, pp. 377–387. ISSN: 02765624. DOI: [10.1016/j.rssm.2012.08.001](https://doi.org/10.1016/j.rssm.2012.08.001).
- Weidinger, Laura et al. (Dec. 2021). “Ethical and Social Risks of Harm from Language Models”. In: *arXiv:2112.04359 [cs]*. arXiv: [2112.04359 \[cs\]](https://arxiv.org/abs/2112.04359).
- Xu, Frank F., Bogdan Vasilescu, and Graham Neubig (Sept. 2021). “In-IDE Code Generation from Natural Language: Promise and Challenges”. In: *arXiv:2101.11149 [cs]*. arXiv: [2101.11149 \[cs\]](https://arxiv.org/abs/2101.11149).
- Young, C. (Dec. 2012). “Losing a Job: The Nonpecuniary Cost of Unemployment in the United States”. In: *Social Forces* 91.2, pp. 609–634. ISSN: 0037-7732, 1534-7605. DOI: [10.1093/sf/sos071](https://doi.org/10.1093/sf/sos071).