

NBER WORKING PAPER SERIES

ENHANCING WORKER PRODUCTIVITY WITHOUT AUTOMATING TASKS:
A DIFFERENT APPROACH TO AI AND THE TASK-BASED MODEL

Ajay K. Agrawal
John McHale
Alexander Oettl

Working Paper 34781
<http://www.nber.org/papers/w34781>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2026

We gratefully acknowledge financial support from the Center for Innovation and Entrepreneurship at the University of Toronto's Rotman School of Management. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w34781>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2026 by Ajay K. Agrawal, John McHale, and Alexander Oettl. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Enhancing Worker Productivity Without Automating Tasks: A Different Approach to AI and the Task-Based Model

Ajay K. Agrawal, John McHale, and Alexander Oettl

NBER Working Paper No. 34781

January 2026

JEL No. J24, O33, O41

ABSTRACT

The task-based approach has become the dominant framework for studying the labor-market effects of artificial intelligence (AI), typically emphasizing the replacement of human workers by machines. Motivated by growing empirical evidence that contemporary AI is more often used as a tool that augments workers, this paper develops two related task-based models in which AI enhances worker productivity without automating tasks. Abstracting from capital, we develop a pair of related task-based models that examine how technological progress in AI that provides new tools to augment workers affects aggregate productivity and wage inequality. Both models emphasize the role of human capital in intermediating the effects of AI-related technological shocks. In the first model, AI use requires specialized expertise, and technological progress expands the set of tasks for which such expertise is effective. We show that a larger supply of AI expertise amplifies the productivity gains from improvements in AI technology while attenuating its adverse effects on wage inequality. The second model focuses on non-AI skills, allowing AI tools to alter the set of tasks that workers can perform given their skills. In equilibrium, workers allocate across tasks in response to wages, generating an endogenous distribution of skills across the task space. A central result is that aggregate productivity and wage inequality depend on different global properties of this equilibrium distribution: productivity is particularly sensitive to thinly staffed tasks that create bottlenecks, while wage inequality is driven by the concentration of workers in a narrow set of tasks. As a result, improvements in AI tools can induce non-monotonic co-movement between productivity and inequality. By linking these mechanisms to multidimensional human capital---including AI expertise and higher-order non-AI skills---the paper highlights the role of education and training policies in shaping the economic consequences of AI-driven technological change.

Ajay K. Agrawal
University of Toronto
Rotman School of Management
and NBER
ajay.agrawal@rotman.utoronto.ca

Alexander Oettl
Georgia Institute of Technology
and NBER
alexander.oettl@scheller.gatech.edu

John McHale
University of Galway
john.mchale@universityofgalway.ie

1 Introduction

The task-based model with workers and capital has become the most influential model for studying the economics of artificial intelligence (AI) (Acemoglu and Restrepo, 2018, 2019; Acemoglu, 2025; Restrepo, 2025).¹ Although the model is flexible enough to allow for channels through which AI can augment workers,² the focus is typically on automation—the replacement of workers by machines. This is consistent with the emphasis in the AI engineering literature on developing AI that can match or exceed human workers in the performance of *all* cognitive tasks, often driven by the goal of developing artificial general intelligence (AGI).³

A strength of task-based models compared to earlier skill-biased technical change models is that it can explain how a fall in absolute wages could occur in response to certain types of technological change (Acemoglu and Autor, 2011). It can also explain the effects of advances in AI technology on aggregate productivity. This combination has sometimes led to relatively pessimistic conclusions, where the wage distribution can be severely disrupted but there are only “so-so” effects on productivity (Acemoglu and Restrepo, 2019; Acemoglu, 2025). One reason for these pessimistic conclusions is that standard versions of the model tend to ignore the possible effects of AI on the economy’s knowledge (or ideas) production function. However, another strand of the task-based literature has incorporated tasks directly into the knowledge production function (Aghion et al., 2019; Jones, 2025). By changing the economy’s growth rate through increased scientific output and innovation, these models allow AI to have transformative effects on the size of the economy over time (Agrawal et al., 2026).

We make four contributions to the task-based literature on artificial intelligence. First, we develop two tractable task-based models in which AI operates purely as a worker-augmenting tool

¹Important early contributions to the development of the task-based model include Zeira (1998), Acemoglu and Autor (2011) and Autor et al. (2003).

²The main channel is the creation of new tasks for human workers to do.

³While the replacement of workers by machines is the standard channel through which technological development takes place in these models, broader discussions treat the direction of technological change—“replacement versus tool” or “automation versus augmentation”—as at least partially a policy choice. See, for example, Brynjolfsson (2022) and Acemoglu and Johnson (2023). Sendhil Mullainathan has argued that the use of benchmarks reflecting the “common task framework” biases AI development in a replacement direction. He argues instead for an approach that he describes as building “bicycles for the mind” that is focused on developing AI tools that enhance human workers rather than replacing them. (See the January 21, 2025 session on AI and Opportunity organized by the MIT Shaping the Future of Work Initiative for an excellent overview of the issues. Available to watch at: <https://www.youtube.com/watch?v=XXRt-nm4JQQ>.) While we agree that whether AI develops as a replacement for workers or as a tool to support their productivity is partly a design choice, our assumption in this paper is that humans remain essential across tasks with AI serving as a tool; albeit a tool that can affect the skill requirements for tasks.

rather than as a substitute for labor, abstracting from capital to isolate augmentation mechanisms. Second, we show that improvements in AI affect aggregate productivity and wage inequality through two conceptually distinct channels: changes in task-specific productivity parameters and changes in the allocation of workers across tasks induced by altered skill requirements. Third, in a general equilibrium setting with free worker mobility, we demonstrate that aggregate productivity and wage inequality depend on different functionals of the equilibrium skill-allocation schedule—geometric mean uniformity for productivity and Shannon differential entropy for inequality—implying that the two need not move monotonically with one another following AI-induced reallocations. Finally, we connect these theoretical results to the modern human capital literature, highlighting how investments in general skills, task-specific expertise, and AI-related skills differentially mediate the productivity and distributional consequences of AI-driven technological change.

Although we are of course early in the development of AI, the extant empirical evidence does not suggest AI is leading to a large-scale replacement of workers by machines in either output or knowledge production. Instead, the evidence seems more consistent with AI as a productivity-augmenting *tool* used by workers. One reason might be that human judgment remains essential in most tasks—at least for now.

In a large-scale study of ChatGPT use, Chatterji et al. (2025) observe that workers use the LLM for broadly defined “decision-support,” which they find to be especially important in knowledge-intensive jobs. Brynjolfsson et al. (2025) also find evidence of AI’s use in decision-support by customer-support agents. They find that access to AI assistance increases worker productivity, as measured by issues resolved per hour, by 15% on average, but with substantial heterogeneity across workers. The largest benefits appear to go to lower-paid workers. Less experienced and lower-skilled workers show clear improvements in the speed and quality of their output, while the most experienced and highest-skilled workers see only small gains in speed and actually suffer small declines in quality. In a study of management consultants engaging in complex knowledge work, Dell’Acqua et al. (2023) find that consultants using AI for tasks that were “within the frontier of AI capabilities” were significantly more productive and produced significantly higher quality results, compared to a control group. While consultants across the skills distribution benefited from having access to AI tools, those below the average performance threshold gained most. However, for a task selected to be outside the AI capability frontier, they observe that the use of AI actually lowered

performance. In another recent study, Noy and Zhang (2023) examine the productivity effects from ChatGPT in midlevel professional writing tasks, finding that the use of the technology substantially lowered the time taken to complete assignments and increased the quality of the output. Finally, Peng et al. (2023) present results from a controlled experiment with GitHub Copilot, an “AI pair programmer.” The treatment group completed the experimental task 55.8 percent faster than the control group. They find that developers with less programming experience, older programmers, and those who program more hours per day had the largest productivity gains, concluding that the results point to the promise of AI-programmer pairs to expand access to careers in software development.

The choice between a replacement-centered view and a tool-centered view matters for determining an appropriate policy response to transformative AI. The replacement-centered view naturally focuses attention on major challenges such as the need to develop an alternative income distribution system (e.g., universal basic income (UBI)). Human capital-focused responses would be largely futile if AI is going to replace most workers anyway. However, human-capital investment policies take on greater importance under a tool-centered view. Ensuring that workers are equipped with the expertise to use AI tools and have the necessary complementary skills in a world where AI tools are ubiquitous is likely to be a critical part of the policy response.

Motivated by the current real-world dominance of the AI-as-tool paradigm, we develop a pair of related task-based models in which AI is purely worker augmenting and explore their human-capital policy implications. We abstract from capital in the models to focus more cleanly on tool-based mechanisms that interest us. Thus, we treat AI as (to a first approximation) freely available tools that augment individual-worker productivity.⁴ Both models assume the same underlying task-based production function. We take the task structure as given but allow the availability of AI tools to change the skill requirements for any given task.⁵ To allow for transformative growth effects, we

⁴Completely abstracting from capital is obviously a drastic simplification, particularly given the massive investments we observe taking place to build the capacity for model training and inference. However, despite these large investment costs, the costs for many users for more standard models are often zero or involve a relatively low monthly cost. The availability of open source models, and the possible evolution of ad-based business models, may also be forces keeping the direct costs to users low. Of course, these factors may change in the future depending on how competitive dynamics among the main industry players evolve. Clearly, however, the direct costs of more specialized uses of AI are not zero. We thus treat AI technology as effectively a public good for modelling convenience. It allows us to focus on the importance of human skills—both expertise in the use of AI itself and the importance of complementary non-AI skills—in relatively tractable settings.

⁵In any task-based model, there are important decisions to make in how to conceptualize tasks and skills. In particular, technological change might be captured as changes in the tasks or as changes in the skills required to

take this production function to be the economy’s knowledge production function. However, all the results can be easily reinterpreted in terms of the economy’s final-good production function.

In Model 1, there are two types of workers: AI-expert workers and ordinary workers; and two types of tasks: tasks that are conducive to using AI as a tool and tasks that are not. Only AI-expert workers can use AI in an AI-conducive task. Technological improvements in AI take place at the extensive margin (i.e., the share of tasks for which it is feasible to use AI as a tool). Thus, the advancement of AI is associated with a replacement effect in certain tasks—but it is the replacement of ordinary workers with AI-expert workers not the replacement of workers by machines. We examine the effects of improvements in AI on both aggregate (knowledge) productivity and wage inequality. Just as in the standard task-based model, technological improvement in AI will raise both aggregate productivity and wage inequality. However, a key result is that the share of AI-expert workers in the workforce (our measure of human capital in Model 1) amplifies the impact of a technological improvement on productivity while also mitigating its adverse impact on wage inequality.

While Model 1 provides a transparent setting to examine the importance of AI expertise, it completely abstracts from heterogeneity in workers’ non-AI-related skills. Model 2 introduces a different mechanism through which AI tools can affect the labor market. Now, each task has its own skill requirements, but these requirements can be affected by the availability of AI tools (see also Autor, 2024; Autor and Thompson, 2025). Improvements in AI can open up task opportunities by reducing the need for specific skills or close off tasks where the AI tool is highly complementary to specific skills. For each worker, the state of AI technology determines the set of task markets in which the worker can compete. The actual distribution of workers across tasks is then determined in the general equilibrium of the model. Improvements in AI can open up task opportunities by reducing the need for specific skills or close off tasks where the AI tool is highly complementary to particular skills.

In this second model, we find that both aggregate (knowledge) productivity and wage inequality

complete a task. For example, a central feature of Acemoglu and Restrepo (2018, 2019) is that technological change can create new tasks. We choose to treat tasks as fixed and to model technological change as changing the skills required across the given space of tasks. Although we treat such advancements in AI tools as worker-augmenting, an alternative approach is to view a given task (or occupation) as a bundle of sub-tasks. The availability of a better AI tool might then be seen as automating certain sub-tasks. This is the approach taken by Autor and Thompson (2025). For our modelling, we find it more straightforward to think of an improved AI tool as directly augmenting a worker in a given task. However, we share with Autor and Thompson an emphasis on how access to the AI tool can change the *skill requirements* for the task.

depend on the uniformity of skill densities over tasks in the equilibrium of the model. Interestingly, productivity and wage inequality depend on different measures of uniformity. Aggregate productivity depends on an index of Mean Geometric Uniformity of the equilibrium skill densities; however, wage inequality depends (inversely) on the Shannon Differential Entropy of the equilibrium skill densities. While productivity and wage inequality will typically move inversely in response to an improvement in the AI technology, we show this is not always the case. Productivity is particularly sensitive to “bottlenecks” in the skill densities; in contrast, wage inequality is sensitive to “concentrations” in the skill densities. The model thus provides a relatively rich menu of possibilities for the aggregate implications of improvements in AI tools. How the economy will actually respond will depend on the details of the actual distribution of human capital across workers as well as how AI affects the skill requirements of tasks.

For both models, we motivate our treatment of AI-as-a-tool with the assumption that human workers have continuing roles across tasks because of the enduring importance of human judgment. This view overlaps with the emphasis on “higher-order skills” in the modern human capital literature (see Deming, 2022; Deming and Silliman, 2025, for recent overviews). We draw on this literature in our discussion of the possible implications of the models for human capital development policies in the age of AI.

We organize the rest of the paper as follows. In Section 2, to motivate our focus on AI as a tool, we first consider why human intelligence might remain robustly important across tasks even in the presence of rapidly improving AI. In Section 3, the core of the paper, we set out our two related task-based models of AI as a tool for augmenting workers. We discuss the implications of the results for human capital policies in Section 4 and conclude in Section 5.

2 Why Might Human Skills Remain Robustly Important Across Tasks?

Our working assumptions are that some level of human skill remains essential in all tasks and that AI provides valuable tools to assist workers in those tasks (and indeed will change the nature of the skills required for many of those tasks). For the former assumption to be plausible there must be some robustly important features of human intelligence that are hard to replicate with AI. While

there are obvious risks to being overly prescriptive about what these features are, for the sake of concreteness we provisionally locate them in distinctive human capabilities to make judgments of causal and normative significance about states of the world and to make creative judgments in constructing models of the world. Agrawal et al. (2018b,a, 2019b,a, 2022, 2025) explicitly highlight the role of human judgment relative to AI-based predictions.

We think of conscious human agents as having both a subjective point of view and a (more-or-less integrated⁶) model of the objective physical/social world (or “world model”) that includes agents with subjective viewpoints within it.⁷ The world model—much of which may be tacit⁸—is used by the agent to predict the effects of actions or events on states of the world, including ultimately the effects on own welfare, others’ welfare and the non-welfare related goals of the agent.⁹ We take it that this interrelationship between the subjective and objective viewpoints is what allows human agents to see the causal and normative *significance* of states of affairs, including the significance any changes in states of affairs that result from their actions.¹⁰ Taken as a set of predictive and

⁶The integrated nature of the agent’s world model (or “web of beliefs”) seems to be central to the flexible way that human agents can act across different contexts. The idea of “core knowledge” may help explain how this integration is achieved. While the idea of innate core knowledge is controversial, it is sufficient to assume that some “ideas” are innate or are acquired early in child development. Ideas such as objecthood, spatio-temporal location, causality and the belief/desire psychology of other people seem fundamental in the construction of an integrated model of the world.

⁷See Nagel (1986) for a fascinating philosophical exploration of interrelationships between the subjective and objective points of view. Among philosophers, John Searle is one of the most insistent on the importance of the subjective viewpoint, arguing that, while this viewpoint is causally reducible to physical processes in the brain, it is “ontologically irreducible” (see, e.g., Searle, 2004). While Searle is dogmatic that consciousness is a feature only of biological brains and argues against the possibility of what he called “strong AI,” we do not share that dogmatism. Instead, while allowing for the theoretical possibility of AGI, we simply assume that for now we remain some distance from achieving it.

⁸Following Polanyi, tacit knowledge is often formulated as “knowing more than you can tell” (see, for example, Polanyi, 1966). But it also seems that in communication we often “tell more than we know”. This can take the form of using concepts for which we have only a vague sense of their meaning, implicitly appealing to the authority of experts—saying essentially that we mean whatever it is the experts mean. One way of viewing LLMs is that they are adept at telling more than they know; indeed, on a reasonable view of knowledge, they can be seen as knowing nothing at all in terms of having a sense of the significance in the world of what they are communicating. Even so, they can be remarkable sources of information and aids to reasoning. However, their inability to judge the real world significance of what they are communicating may mean that for many tasks the human agent must remain in the decision-making loop.

⁹See Sen (1987) for an argument for the importance of integrating ethical concerns into the study of economic behavior.

¹⁰In economics, we typically assume a belief-desire psychology that makes a sharp distinction between facts (the objective model underlying beliefs) and values (as determined subjectively by the agent). Many philosophers see a fuzzier line between facts and values (see, e.g., Putnam, 2002); in particular, values are allowed to be partly determined by the reasoning agent—reasoning that involves both the use and creation of their world model. This opens up the possibilities for more complex ethical and social behavior than we typically allow for in our economic models but may be important in the accomplishment of many tasks. Our formulation is sufficiently flexible to allow for both a narrow and wide treatment of the fact/value distinction. Even on the narrower treatment, AI would face a significant challenge in replicating the behavior of a human agent as they effectively would have to know their utility function. The challenge becomes even greater when behavior is the result of a complex mix of self-interested, social

generative tools, AI may assist humans in performing any given task, but we take it that this human capacity to make judgments of significance will remain robustly important, at least over the medium term.¹¹

A second source of enduring differential human capability may lie in the fluidity with which human agents can continuously construct and adapt their world models, often with relatively limited data.¹² Because they possess an integrated world model, humans are capable of being *surprised* by the presence or absence of things in the world, leading them to ask substantive questions. Moreover, their integrated models allow them to search through vast combinatorial spaces of existing knowledge to form hypotheses about the answers to those questions. These *creative* capabilities have received recent attention in relation to the difficulties that AI has compared to humans in completing puzzles in the ARC Challenge (see, e.g., Chollet, 2019). Solving these puzzles are archetypical examples of Peircean abductive reasoning.¹³ Based on few examples, the reasoner must draw on their existing stock of knowledge to form a hypothesis about the simple algorithm that solves the puzzle, which is then tested on one or more held-out examples. Such hypothesis formation through abductive “guessing” followed by inductive testing (and/or argument) seems to be central to the development of world models, whether it is the quotidian models used to navigate the day-to-day world or the more formal scientific varieties.

Given the vast combinatorial space of assumptions that could underlie any given model, it remains mysterious how humans can use their existing (integrated) knowledge stocks to identify plausible hypotheses. The answer may lie in the interrelationship between Kahneman’s intuitive System 1 in narrowing the search space combined with the deliberative powers of System 2 in at

and ethical concerns.

¹¹The need to move beyond language to develop more complete “world models” is increasingly being recognized by AI developers and researchers. One major focus is on the use of video data to predict events in the physical world, with impressive results being achieved in terms of realistic video generation by products such as Sora (OpenAI) and Veo (Google). Google’s Genie also allows the user to interact with the program in generating the video. Fei-Fei Li’s World Labs is also exploring “spatial intelligence,” including methods to overcome the challenges of “persistence” with existing models. Yann LeCun has argued against a focus on generative models to produce realistic video, instead focusing on models such as JEPA that support planning based on abstract physical models of the world (see LeCun, 2022). For all these approaches, it is unclear whether these statistical models can “understand” causal structure in the world in the way that humans do. One hope is that causal understanding will emerge with some combination of the scaling of these models and algorithmic improvements. However, the models seem to be far from human-level understanding of the physical world and nowhere close to human-level understanding of the social world.

¹²See Silver and Sutton (2025) on potential routes to continuous learning in AI.

¹³As set out by the American pragmatist philosopher Charles Sander Peirce, the canonical form of abductive reasoning has the structure: The surprising fact, *C*, is observed; But if *A* were true, *C* would be a matter of course; Hence there is reason to suppose that *A* is true (Peirce, 1998).

least provisionally assessing candidate options presented by System 1 (Kahneman, 2011). While it is certainly conceivable that AI (or a set of partly task-specific AIs) could acquire human-like model-construction capabilities, at present it appears to struggle to match the fluidity of human judgments of hypothesis plausibility across a wide range of domains.¹⁴

While our models take as given the essential nature of human workers in the completion of tasks, it also allows AI to provide powerful tools (or even to serve as “partners”) to aid with those tasks. In the context of knowledge production, we find it useful to distinguish between knowledge search across the *idea space* and search across the *design space*. AI can provide powerful tools for each (see Agrawal et al., 2019, 2023, 2024, 2026). The idea space is a (potentially vast) space of possible combinations of existing ideas, where a successful new idea is some previously undiscovered combination. LLMs, for example, provide tools to aid this search, including for locating existing ideas and for exploring how those ideas might be combined. Anyone who has used the latest generations of LLMs is likely to come away at least somewhat impressed by its creative combinatorial powers in response to a human-supplied prompt, notwithstanding its occasional confabulations.

In contrast to the idea space, we think of a design space as a physical space of combinations—say the amino acid sequences that determine a protein’s structure. Such spaces are notoriously difficult for a human scientist to search. Predictive and generative AI models (such as AlphaFold in the case of protein shape prediction) are already proving valuable aids to search. Even so, there remain roles for human judgment in identifying the important questions and assessing the plausibility of model outputs before sending for expensive testing. While there may be cases where it makes most sense to think of AI as completely automating some task (i.e., the machine replacing the worker), our working assumption will be that workers retain at least *some* role in *all* tasks. This may be the case even if the requirements for human expertise in the task are dramatically reduced, potentially opening up the task to a much larger number of workers.

We have focused to this point on the distinctive capability of human workers to make judgments of significance, including the causal significance of events or actions and the normative significance of

¹⁴There is of course active debate among AI experts as to the most plausible route to human-level intelligence (or AGI). Some believe that the scaling (along with algorithmic improvements) of LLMs will eventually lead to AGI. Others see the need for different architectures to match the type of abductive reasoning capabilities evident in humans (e.g., probabilistic-program-induction-type architectures (Lake et al., 2017)). More recently, there has been significant interest in neuro-symbolic methods that combine the strengths of connectionist and symbolic architectures. One major challenge in matching the human capability to see significance is how to “ground” the AI’s “knowledge” in the real world. (See Harnad, 1990, for a classic exposition of the symbol grounding problem.)

resulting states of affairs given their own or others’ objectives. We have also noted that the creative process of building world models is itself a judgment-laden process. Of course, there is variation in the quality of the world models that workers use to make judgments—e.g., there is variation in the quality of causal models that workers use to predict the outcomes of interventions they might make in the world. We thus next connect the quality of world models that workers are deploying to the human capital (or skills) a worker possesses. Our key assumption is that AI tools can affect the skills required to complete a given task as well as their (AI-augmented) productivity in that task. So, while it is taken as given that human workers will remain necessary even with advanced AI, the nature of the skills that are required in the completion of any given task can change as AI develops.

The emphasis on the capability to make good judgments has notable overlaps with the modern approach to human capital (e.g., Heckman and Kautz, 2012; Deming, 2017, 2022; Deming and Silliman, 2025). Compared to the traditional human capital literature that typically concentrated on single dimensional measures of cognitive human capital such as IQ or years of schooling, the modern approach explicitly treats human capital as a multidimensional construct, including “non-cognitive” characteristics such as perseverance, attentiveness, motivation and self-confidence (Heckman, 2013).¹⁵ There has also been a stress on the “technology of skill formation,” including the importance of dynamic complementarities between earlier and later investments in different types of skill (Cunha and Heckman, 2007).

The modern approach has emphasized success determinants that were largely absent from the classic human capital literature, such as self-regulation or “character” skills (James Heckman) and social skills in team settings (David Deming).¹⁶ Heckman’s stress is mainly on what might be

¹⁵Although Heckman labels such characteristics as non-cognitive—which we interpret as behaving according to instinct or habit—we do not see why they cannot (at least sometimes) be a form of cognitive ability to see reasons for acting in particular ways.

¹⁶The modern approach to human capital can be seen as blurring the fact/value distinction that is typically taken as paradigmatic in economic models. For example, the ability to delay gratification is seen as a “skill” that can be developed rather than simply a reflection of exogenously given preferences. As another example, a willingness to reciprocate is partly a “skill” that can be developed rather than simply as a feature of the equilibrium in a repeated game played by purely self-interested players. If human economic agents were really sufficiently well captured by models that assume a strict separation between beliefs and preferences (or between facts and values), the potential for AI agents taking over from human agents in most tasks would probably be high. It is human agents’ ability to see significance in a complex world that makes them hard to replace. Although we allow for a wider range of motivations than might be typical in an economic model, our focus is still on an individual’s success in the labor market. However, this success now depends on the individual’s ability to make complex judgments of significance in the world—judgments that we assume are hard or impossible for an AI. This differential capacity underlies our assumption that humans remain essential in all tasks, albeit potentially significantly supported by AI tools in many tasks.

called “intrapersonal judgment”—e.g., the ability to achieve longer-term goals by making appropriate trade-offs between longer-term success and instant gratification. In contrast, in the context of social interaction, Deming emphasizes what might be called “interpersonal judgment”—e.g., being able to notice social cues from partners in a team. Both can reflect an ability to judge *reasons* for acting in an appropriate way in a given context. Although perhaps drawing too sharp a line between cognitive and non-cognitive skills, Heckman and Corbin (2016) argue for the importance of looking beyond more narrowly defined cognitive skills (say as measured by IQ) to a broader set of skills that can affect both of these forms of judgment:

An important lesson from the recent economics of human development is that cognitive skills are only part of what is required for success in life. Personality skills—i.e., “soft skills,” like trust, altruism, reciprocity, perseverance, attention, motivation, self-confidence, and personal health—are also important. Health and nutrition—macro- and micro-nutrients—are essential skills. So are the abilities to make wise decisions, to guide one’s life by reflective reason, and to plan ahead. These skills are often neglected in scientific analyses and policy discussions alike. (Heckman and Corbin, 2016, p. 10)

We recognize that there may be cases where an AI’s lack of bias might make it superior to human workers—for one thing, AI does not face the same Heckman-type challenges of self-regulation. It is noteworthy that Daniel Kahneman expressed sufficient frustration with human foibles—including the noisiness of much human decision-making—that he saw the case for AI taking over many tasks:

[M]ost of the errors people makes are better viewed as random noise, and there is an awful lot of it. Admitting the existence of noise has implications in practice. One implication is obvious. You should replace humans for algorithms whenever possible. (Kahneman, 2019, p. 610)

But this quote catches Kahneman in a particularly pessimistic mood about human capabilities. He was also deeply aware of the positive powers of human judgment. Here is what he had to say in the concluding chapter of *Thinking, Fast and Slow*:

I have spent more time describing System 1, and have devoted many pages to the errors of intuitive judgment and choice that I attribute to it. However, the relative number of pages is a poor indicator of the balance between the marvels and flaws of intuitive thinking. System 1 is indeed the origin of much that we do wrong, but it is also the origin of much of what we do right—which is most of what we do. Our thoughts and actions are routinely guided by System 1 and are generally on the mark. One of the marvels

is the rich and detailed model of our world that is maintained in associative memory; it distinguishes normal from surprising events in a fraction of a second, immediately generates an idea of what was expected instead of a surprise, and automatically searches for some causal interpretation of surprises and events as they take place. (Kahneman, 2011, pp. 415–416)

This is the worker’s rich, complex and human-capital underpinned world model—or at least the more intuition-based aspects of that model—in action.

However, the boundary between human and machine capabilities is dynamic, particularly where judgment can be explicitly defined. Judgment can be codified when we can articulate in advance the values we ascribe to different outcomes or “states of the world”. When we can codify judgment in this way, we can fully automate decisions or actions, which are based on the predictions of the likelihood of different outcomes and the judgment of the relative value of each possible outcome. Perhaps one of the most valuable recent inventions in AI was the ability to codify judgment in language models through a process called Reinforcement Learning with Human Feedback (RLHF). This process works by having human evaluators rank model outputs to create a “reward model,” which effectively serves as a digital proxy for human judgment regarding the quality and safety of a response. The AI then optimizes its policy against this codified reward signal, learning to generate output that aligns with human preferences rather than simply predicting the next statistically probable token. This innovation of RLHF unlocked Large Language Models, transforming them from interesting but unreliable prediction engines into useful, instruction-following tools that can act in accordance with human intent.

Overall, we stress that our assumption is not that AI systems are incapable in principle of performing all aspects of judgment, but rather that currently deployed and foreseeable AI tools function primarily as decision-support systems whose effectiveness depends critically on human interpretation, goal setting, and evaluation. In this sense, the continued role of human workers in our models reflects both current empirical evidence and a modeling choice intended to isolate augmentation mechanisms rather than to forecast the ultimate limits of artificial intelligence. Therefore, while allowing for human judgment to be enhanced by access to AI tools, in developing our task-based models we simply take it as given that human judgment is essential to some degree in all tasks. We then examine how the effects of improved AI tools on productivity and wage inequality are intermediated by AI-related (Model 1) and non-AI-related (Model 2) human skills. We then discuss

some possible implications for human capital policy in the age of AI.

3 The Effects of Improvements in AI on Aggregate Productivity and Wage Inequality

For both models, we assume a constant-returns-to-scale Cobb–Douglas aggregator over tasks. This assumption is not intended to restrict substitution possibilities across tasks in a realistic sense, but rather to isolate and clarify the reallocation mechanisms that are central to our analysis. The Cobb–Douglas structure ensures that relative wages depend only on the equilibrium distribution of workers across tasks and not directly on task-specific productivity parameters. This separation allows us to distinguish sharply between the effects of AI on aggregate productivity—through changes in task productivity and task coverage—and its effects on wage inequality, which operate through worker reallocation across tasks. More general aggregators would introduce additional channels through which task-specific productivity affects relative wages, but would not overturn the core mechanisms emphasized here. We therefore view the Cobb–Douglas case as a useful benchmark that highlights how global properties of the equilibrium allocation of skills across tasks shape productivity and inequality in a task-based economy.

3.1 Model 1: The Importance of Skills in the *Use* of AI

Our first model focuses simply on how the availability of AI expertise intermediates the effects of an improvement in the AI technology on aggregate knowledge productivity and wage inequality. To focus clearly on the effects of the supply of AI expertise, we assume there are just two types of workers: AI-expert workers and ordinary workers. We also assume that there are tasks for which it is feasible to use AI expertise and tasks for which it is not. (We relax these assumptions in Model 2.) We adapt the Acemoglu-Restrepo task-based model to this pure augmentation scenario (see Acemoglu and Restrepo, 2019, for an accessible introduction to the model for the case of a Cobb–Douglas aggregator). As in Acemoglu and Restrepo (2019), we assume a constrained comparative advantage equilibrium—i.e., at the current equilibrium wages, firms would employ AI expertise for more tasks if it were technologically feasible to do so.

We focus on technological improvements in AI on the *extensive* margin. Starting from a con-

strained equilibrium, we examine the effects of a marginal increase in the measure of tasks that it is feasible to complete using AI-expert workers. Our main result relates to how the relative supply of AI expert workers intermediates the effect of the extensive-margin technological improvement on aggregate productivity and wage inequality.

More formally, we assume there is a continuum of tasks required for knowledge production. These tasks are arranged on $[0, 1]$. For a given task, τ , the output of the task is given by $z(\tau)$. We denote existing stock of knowledge as A and the output of new knowledge as $\dot{A}$. As noted, the production function (or aggregator) for knowledge takes the constant-returns-to-scale Cobb-Douglas form:

$$\dot{A} = A \left(\exp \left[\int_0^1 \ln z(\tau) d\tau \right] \right). \quad (1)$$

The skill density in task τ is denoted $s(\tau)$.¹⁷ We assume there is a continuum of workers distributed across the task distribution such that $\int_0^1 s(\tau) d\tau = \int_0^1 di = 1$, for $\tau \in [0, 1]$ and $i \in [0, 1]$, where i indexes a worker.¹⁸ There are two types of workers: ordinary workers without AI expertise (which we index by O); and workers with AI expertise (which we index by AI). The shares of ordinary and AI-expert workers are s_O and s_{AI} respectively. The density of ordinary workers in a given task, τ , is given by $s_O(\tau)$ and the density of AI-expert workers in that task is $s_{AI}(\tau)$, for $\tau \in [0, 1]$. Given the state of AI technology, the measure of tasks that utilize AI expertise is $\tau \in [0, \bar{\tau}]$. All other tasks, $\tau \in (\bar{\tau}, 1]$, could be done by ordinary or AI-expert workers. We assume $0 < \bar{\tau} < 1$. (Note that we also assume there is free disposal of AI expertise, so that AI-expert workers are free to work as ordinary workers. However, we later impose conditions on the model such that no worker with AI expertise has an incentive to work as an ordinary worker.)

The productivity of an ordinary worker in any given task τ is $\theta_O(\tau)$; the productivity of an AI-expert worker in an AI-supported task is $\theta_{AI}(\tau)$, where we refer to the θ s as task-specific productivity parameters. To rule out degenerate output, we assume that $\theta_O(\tau), \theta_{AI}(\tau) > 0$ almost everywhere and that $\int_0^1 \ln \theta(\tau) d\tau > -\infty$.¹⁹ A key assumption of the model is that $\theta_O(\tau)/\theta_{AI}(\tau)$ is increasing in τ , and so ordinary workers have a *comparative* advantage in higher indexed tasks.

¹⁷Throughout, we use the term “skill density” as shorthand for the density of workers supplying their skills to a given task in equilibrium; skills do not exist independently of workers.

¹⁸As usual in task-based models where tasks are defined on a continuum, tasks are indexed by the unit interval equipped with the Borel σ -algebra and Lebesgue measure.

¹⁹We also assume that the task-specific productivity parameters are bounded from above: i.e., $\theta(\tau) \leq \bar{\theta} < \infty$, for all $\tau \in [0, 1]$.

The share of tasks that it is technologically feasible to perform using AI-expert workers is $\bar{\tau}$, so that a change in the share of tasks that can be done with AI expertise is the source of (extensive-margin) AI-related technological change in the model. The current state of technology therefore provides a constraint on the tasks that are feasible for AI-expert workers; a constraint we assume is binding (see Acemoglu and Restrepo, 2019, for a parallel assumption in the case where machines can replace workers). In other words, at the current allocation of tasks to ordinary and AI-expert workers, AI-expert workers would be used for the next task in the index *if* it were technologically feasible. The technological constraint thus prevents the optimal allocation of workers to tasks based on comparative advantage leading to a *constrained* comparative advantage equilibrium.

This constraint can be expressed as:

$$\frac{w_O}{\theta_O(\bar{\tau})} > \frac{w_{AI}}{\theta_{AI}(\bar{\tau})}. \quad (2)$$

That is, the unit cost of producing using ordinary workers is greater than the unit cost of producing using AI expert workers at $\bar{\tau}$. We capture AI-related technological progress at the extensive margin as an exogenous marginal increase in $\bar{\tau}$.

With these constraints, and with complete specialization, the output of any given task, τ , is:

$$z(\tau) = \begin{cases} \theta_{AI}(\tau)s_{AI}(\tau) & \text{if } \tau \in [0, \bar{\tau}] \\ \theta_O(\tau)s_O(\tau) & \text{if } \tau \in (\bar{\tau}, 1] \end{cases}. \quad (3)$$

Given competitive task markets, the price per unit of a task is equal to marginal/average cost of that unit. This cost is equal to the wage divided by the productivity of the relevant worker in the task:

$$p(\tau) = \begin{cases} \frac{w_{AI}}{\theta_{AI}(\tau)} & \text{if } \tau \in [0, \bar{\tau}] \\ \frac{w_O}{\theta_O(\tau)} & \text{if } \tau \in (\bar{\tau}, 1] \end{cases}. \quad (4)$$

Normalizing the price of a unit of final output to unity and assuming competitive skills markets, the quantity demanded of any task is equal to output divided by the price of the task:

$$z(\tau) = \frac{\dot{A}}{p(\tau)}. \quad (5)$$

The quantity of AI expert workers demanded in each task is then:

$$\frac{z(\tau)}{\theta_{AI}(\tau)} = \begin{cases} \frac{\dot{A}}{w_{AI}} & \text{if } \tau \in [0, \bar{\tau}] \\ 0 & \text{if } \tau \in (\bar{\tau}, 1] \end{cases}. \quad (6)$$

And the quantity of ordinary workers demanded in each task is:

$$\frac{z(\tau)}{\theta_O(\tau)} = \begin{cases} 0 & \text{if } \tau \in [0, \bar{\tau}] \\ \frac{\dot{A}}{w_O} & \text{if } \tau \in (\bar{\tau}, 1] \end{cases}. \quad (7)$$

Aggregating over the relevant tasks, the aggregate demand for workers of the two worker types are:

$$\int_0^1 \frac{z(\tau)}{\theta_{AI}(\tau)} d\tau = \bar{\tau} \frac{\dot{A}}{w_{AI}}, \quad \int_0^1 \frac{z(\tau)}{\theta_O(\tau)} d\tau = (1 - \bar{\tau}) \frac{\dot{A}}{w_O}. \quad (8)$$

Equating the aggregate demand for each expertise type with the fixed supplies of workers of that type allows us to solve for the equilibrium expertise-relevant wages:

$$w_{AI} = \bar{\tau} \frac{\dot{A}}{s_{AI}}, \quad w_O = (1 - \bar{\tau}) \frac{\dot{A}}{s_O}. \quad (9)$$

We next examine the effects on both aggregate productivity and wage inequality of an extensive-margin improvement in the AI-technology.

Aggregate (Knowledge) Productivity

Using $z(\tau) = \theta(\tau)s(\tau)$ and recalling that the size of the workforce has been normalized to one, we can write aggregate productivity as:

$$\dot{A} = A \left(\exp \left[\int_0^1 \ln(\theta(\tau)s(\tau)) d\tau \right] \right). \quad (10)$$

Now taking logs, gathering the task-specific productivity terms and assuming full specialization, we obtain:

$$\ln \dot{A} = \ln A + \left[\int_0^{\bar{\tau}} \ln \theta_{AI}(\tau) d\tau + \int_{\bar{\tau}}^1 \ln \theta_O(\tau) d\tau \right] + \left[\int_0^{\bar{\tau}} \ln s(\tau) d\tau + \int_{\bar{\tau}}^1 \ln s(\tau) d\tau \right]. \quad (11)$$

We next use the relevant demand functions for workers in specific tasks to write the log productivity function in equilibrium as:

$$\ln \dot{A} = \ln A + \left[\int_0^{\bar{\tau}} \ln \theta_{AI}(\tau) d\tau + \int_{\bar{\tau}}^1 \ln \theta_O(\tau) d\tau \right] + \left[\int_0^{\bar{\tau}} \ln \left(\frac{\dot{A}}{w_{AI}} \right) d\tau + \int_{\bar{\tau}}^1 \ln \left(\frac{\dot{A}}{w_O} \right) d\tau \right]. \quad (12)$$

Substituting the equilibrium wages from equation (9):

$$\begin{aligned} \ln \dot{A} &= \ln A + \left[\int_0^{\bar{\tau}} \ln \theta_{AI}(\tau) d\tau + \int_{\bar{\tau}}^1 \ln \theta_O(\tau) d\tau \right] \\ &\quad + \left[\int_0^{\bar{\tau}} \ln \left(\frac{\dot{A}}{\bar{\tau} \frac{\dot{A}}{s_{AI}}} \right) d\tau + \int_{\bar{\tau}}^1 \ln \left(\frac{\dot{A}}{(1 - \bar{\tau}) \frac{\dot{A}}{s_O}} \right) d\tau \right] \\ &= \ln A + \left[\int_0^{\bar{\tau}} \ln \theta_{AI}(\tau) d\tau + \int_{\bar{\tau}}^1 \ln \theta_O(\tau) d\tau \right] + \left[\bar{\tau} \ln \left(\frac{s_{AI}}{\bar{\tau}} \right) + (1 - \bar{\tau}) \ln \left(\frac{1 - s_{AI}}{1 - \bar{\tau}} \right) \right]. \end{aligned} \quad (13)$$

Finally, taking the exponential of both sides we obtain an expression for aggregate knowledge productivity as a function of the existing knowledge stock, the task-specific productivities, the state of the AI technology and the shares of the two types of worker:

$$\dot{A} = A \left(\exp \left[\int_0^{\bar{\tau}} \ln (\theta_{AI}(\tau)) d\tau + \int_{\bar{\tau}}^1 \ln (\theta_O(\tau)) d\tau \right] \right) \left(\left(\frac{s_{AI}}{\bar{\tau}} \right)^{\bar{\tau}} \left(\frac{1 - s_{AI}}{1 - \bar{\tau}} \right)^{1 - \bar{\tau}} \right). \quad (14)$$

It can be seen that this equilibrium form of the knowledge production function takes the general form of the Romer ideas production function.

Proposition 1.1: *Aggregate productivity is non-decreasing in the share of AI expertise in the workforce.*

Proof

$$\frac{\partial \ln \dot{A}}{\partial s_{AI}} = \frac{\bar{\tau}}{s_{AI}} - \frac{1 - \bar{\tau}}{1 - s_{AI}} \geq 0 \quad \text{given } \bar{\tau} \geq s_{AI}. \quad (15)$$

Recall that we assume that AI-expert workers are free to work as ordinary workers. This implies that the wage of AI-expert workers should not fall below the level of ordinary workers; i.e.: $w_{AI} \geq w_O$. Given our wage expressions (9) above, this implies that $\bar{\tau} \geq s_{AI}$ in the equilibrium of the model. This condition is sufficient to ensure that aggregate productivity is non-decreasing in the share of AI-expert workers. Aggregate productivity will be strictly increasing with this share

where the wage of AI-expert workers exceeds the wage of ordinary workers, which we assume is the usual case.

Proposition 1.2: *Aggregate productivity is increasing in the share of tasks for which it is feasible to use AI expertise.*

Proof

$$\frac{\partial \ln \dot{A}}{\partial \bar{\tau}} = \ln \theta_{AI}(\bar{\tau}) - \ln \theta_O(\bar{\tau}) + \ln \left(\frac{s_{AI}}{\bar{\tau}} \right) - \ln \left(\frac{1 - s_{AI}}{1 - \bar{\tau}} \right). \quad (16)$$

$$= \ln w_O - \ln \theta_O(\bar{\tau}) - (\ln w_{AI} - \ln \theta_{AI}(\bar{\tau})) > 0. \quad (17)$$

The second equation follows from our assumption of a constrained comparative advantage equilibrium (equation (2)) and the wage equations (9). Improvements in the AI technology at the extensive margin thus relaxes the constraint and results in higher aggregate productivity.

The intuition for this result is more clearly seen with the help of Figure 1. The figure shows the constrained comparative advantage equilibrium. At the equilibrium wages, at the technologically determined marginal task, $\bar{\tau}$, the cost per unit of the task is lower using AI expertise than using ordinary workers. A marginal relaxation of the constraint would therefore result in firms replacing ordinary workers with AI expert workers in the tasks that it now becomes feasible to avail of AI expertise. Moreover, the size of the resulting cost saving (and associated aggregate productivity gain) will be larger the larger is the AI-expertise cost advantage at the initial constrained equilibrium.

We next turn to the central result in terms of the importance of the availability of AI-expert skills in intermediating the effects of improvements in AI on aggregate productivity.

Proposition 1.3: *The size of the positive effect on aggregate productivity of a small increase in the share of tasks that can be done by AI-expert workers is increasing in the share of AI-expert workers in the workforce.*

Proof

$$\frac{\partial^2 \ln(\dot{A})}{\partial \bar{\tau} \partial (s_{AI})} = \frac{1}{s_{AI}(1 - s_{AI})} > 0. \quad (18)$$

The intuition for this result can be seen from Figure 1b. A larger share of AI-expert workers will lower the wage of AI-expert workers and raise the wage of ordinary workers. This increases the size of the cost gap at the initial constrained equilibrium, resulting in a larger productivity gain

when the constraint is relaxed due to an improvement in the AI technology.

Wage Inequality

As previously noted, in comparison with the older labor-market models that incorporated skill-biased technological change, an empirically attractive feature of the task-based framework is that it allows in principle for the wages of a particular class of worker to fall in absolute terms due to a change in technology. Returning to our wage equations (9), assuming fixed shares of AI-expert and ordinary workers we can write the effects on log wages following an extensive-margin improvement in the technology as:

$$\frac{\partial \ln w_{AI}}{\partial \bar{\tau}} = \frac{1}{\bar{\tau}} + \frac{\partial \ln \dot{A}}{\partial \bar{\tau}}, \quad (19)$$

$$\frac{\partial \ln w_O}{\partial \bar{\tau}} = -\frac{1}{1 - \bar{\tau}} + \frac{\partial \ln \dot{A}}{\partial \bar{\tau}}. \quad (20)$$

The effect on the wage of AI-expert workers is strictly positive in our constrained comparative advantage equilibrium. These workers can gain from *both* a substitution effect and an aggregate productivity effect. However, while ordinary workers will also gain from the productivity effect, they also suffer from a “displacement effect” as more tasks enter the purview of AI experts.

Cognizant of the enhanced scope for wage effects in the task-based model, we next examine the effects of an (extensive-margin) technology improvement on wage inequality. Of course, there are numerous ways to measure wage inequality, with different measures displaying varying degrees of sensitivity to different parts of the wage distribution. A tractable wage inequality measure for our model (one that we also employ in our second model below) is the simple ratio of the arithmetic mean wage ($\bar{w}$) to the geometric mean wage (w^g).²⁰ This measure has the desirable properties of being scale invariant (i.e., proportional increases in wages will not affect measured inequality) and meeting the Pigou-Dalton transfer condition that a transfer from a richer worker to poorer worker that does not change their rank in the distribution will lower wage inequality. One notable feature of the measure is that it is quite sensitive to changes in wages at the lower end of the wage distribution through the effect of low wages on the geometric mean wage. In empirical applications it would

²⁰This measure has close relationships to other classic inequality measures. For the Atkinson inequality measure with parameter $\varepsilon = 1$, $\text{Atkinson}(1) = 1 - (w^g/\bar{w})$. Moreover, the Generalized Entropy measure with parameter equal to 0 (also known as the Mean Log Deviation or Theil-L) is given by: $GE(0) = \ln(\bar{w}/w^g)$. Thus, there are positive monotonic relationships between the $\bar{w}/w^g$ measure of inequality and both $\text{Atkinson}(1)$ and $GE(0)$. It follows that any of the three measures will rank the inequality of distributions identically.

have the drawback that the measure goes to infinity as any wage in the distribution goes to zero and is undefined if a wage is exactly zero. However, the latter is not an issue for our model as we assume wages are strictly positive almost everywhere across the task continuum.

Noting that the arithmetic mean wage is simply $\dot{A}$ given that the workforce is normalized to one, and workers are the only inputs in our simple economy, our wage inequality measure is:

$$\frac{\bar{w}}{w^g} = \frac{\dot{A}}{\exp(s_{AI} \ln w_{AI} + (1 - s_{AI}) \ln w_O)}. \quad (21)$$

Substituting our expressions for equilibrium wages:

$$\frac{\bar{w}}{w^g} = \frac{\dot{A}}{\exp\left(s_{AI} \ln\left(\bar{\tau} \frac{\dot{A}}{s_{AI}}\right) + (1 - s_{AI}) \ln\left((1 - \bar{\tau}) \frac{\dot{A}}{1 - s_{AI}}\right)\right)}. \quad (22)$$

And dividing both the numerator and denominator by the arithmetic mean wage:

$$\frac{\bar{w}}{w^g} = \frac{1}{\exp\left(s_{AI} \ln\left(\bar{\tau} \frac{1}{s_{AI}}\right) + (1 - s_{AI}) \ln\left((1 - \bar{\tau}) \frac{1}{1 - s_{AI}}\right)\right)}. \quad (23)$$

$$= \exp\left(s_{AI} \ln\left(\frac{s_{AI}}{\bar{\tau}}\right) + (1 - s_{AI}) \ln\left(\frac{1 - s_{AI}}{1 - \bar{\tau}}\right)\right). \quad (24)$$

Finally taking logs, we obtain a simple measure of log inequality that depends on both the state of the AI technology and the relative shares of ordinary and AI workers in the workforce:

$$\ln \frac{\bar{w}}{w^g} = s_{AI} \ln\left(\frac{s_{AI}}{\bar{\tau}}\right) + (1 - s_{AI}) \ln\left(\frac{1 - s_{AI}}{1 - \bar{\tau}}\right). \quad (25)$$

We can now examine the related partial and cross partial derivatives that we explored above for aggregate productivity. Again, our main interest is in how the availability of AI-expertise intermediates the impact of an extensive-margin technological improvement on the economy, this time on wage inequality.

Proposition 1.4: *Wage inequality is non-increasing in the share of AI expertise in the workforce.*

Proof

$$\frac{\partial \ln \frac{\bar{w}}{w^g}}{\partial s_{AI}} = \ln\left(\frac{s_{AI}}{\bar{\tau}}\right) + 1 - \ln\left(\frac{1 - s_{AI}}{1 - \bar{\tau}}\right) - 1 \leq 0 \quad \text{given } \bar{\tau} \geq s_{AI}. \quad (26)$$

Again the inequality will be strict if $w_{AI} > w_O$ so that an increase in the AI-expert share will lower wage inequality.

Proposition 1.5: *Wage inequality is non-decreasing in the share of tasks for which it is feasible to use AI expertise.*

Proof

$$\frac{\partial \ln \frac{\bar{w}}{w^g}}{\partial \bar{\tau}} = -\frac{s_{AI}}{\bar{\tau}} + \frac{1 - s_{AI}}{1 - \bar{\tau}} \geq 0 \quad \text{given } \bar{\tau} \geq s_{AI}. \quad (27)$$

This inequality will again be strict in the case where w_{AI} is strictly greater than w_O . Assuming this is the case, the extensive margin improvement in the AI technology will worsen wage inequality as the wages of AI-expert workers rise relative to ordinary workers. We therefore see that aggregate productivity and wage inequality move in the same direction—both aggregate productivity and wage inequality *rise*—following the technological improvement.

Earlier we saw that a larger share of AI-expert workers *amplifies* the positive productivity effect; we next show that such a larger share *attenuates* the rise in wage inequality. Thus, a larger share of AI-expert workers improves the trade-off between the enhanced productivity and worsened inequality following the technological improvement.

Proposition 1.6: *A larger share of AI-expert workers in the workforce attenuates the size of the adverse effect on wage inequality of a small increase in the share of tasks that can be done by AI-expert workers.*

Proof

$$\frac{\partial^2 \ln \left(\frac{\bar{w}}{w^g} \right)}{\partial \bar{\tau} \partial (s_{AI})} = \frac{-1}{\bar{\tau}(1 - \bar{\tau})} < 0 \quad \text{given } 0 < \bar{\tau} < 1. \quad (28)$$

Summing up on Model 1, we see that the availability of AI expertise matters for how extensive-margins improvements in AI are transmitted to aggregate productivity and wage inequality. Not surprisingly, we find that when the wages of AI-expert workers exceed the wages of ordinary workers an increased share of AI-expert workers will raise productivity and lower wage inequality. However, an improvement in AI-technology will improve productivity but worsen wage inequality. The most important result is that a larger share of AI-expert workers will enhance the effect of an improvement in AI technology on productivity but also dampen the effect of that improvement on wage inequality. These results point to the importance of investing in the skills needed to use AI when that technology

is becoming relevant to an increasing range tasks. We discuss this further in Section 4.

3.2 Model 2: The Importance of *Non-AI* Skills

Our first model assumes that worker skills vary on just a single dimension: whether they have AI expertise or not. Moreover, all workers of a given expertise type receive the same wage and are fully mobile across tasks that utilize their specific level of AI expertise. The assumptions that workers are identical in terms of non-AI skills and that they are only constrained in their mobility by AI expertise are obvious limitations of this model.

For our second model, we now place non-AI skills at center-stage, and assume that the skills required in any given task are task specific and workers' non-AI related skills are such that they may be only able to compete in a subset of tasks (a subset that can be affected by the state of the AI technology). To focus on the importance of non-AI skills, we now assume that AI is available as a tool for *all* tasks and *all* workers have AI expertise—but crucially workers now differ in their non-AI skills.

An improvement in AI technology now has two key effects. First, it changes the task-specific productivity parameters (i.e., the θ s) for a worker who has the requisite non-AI skills to work in the relevant task. This is the *intensive* margin. Second, for any given worker, it changes the subset of tasks that they could compete in given their non-AI skills and the skill requirements across tasks. We call this the *reallocation* margin. For example, the availability of new AI tools that reduce the complexity of tasks from the viewpoint of the worker could open up new task opportunities for workers with more limited non-AI skills. This second effect opens up an important new channel through which improvements in AI tools can affect aggregate productivity and wages: the mobility of workers across tasks. In common with Model 1, however, workers' human capital plays a central intermediating role between changes in the available technology and the effects on aggregate productivity and wage inequality.

More formally, we represent the state of AI technology, T , as:

1. $\{\theta(\tau)\}$ for all $\tau \in [0, 1]^T$; and
2. $[\Omega(i) \subseteq [0, 1]$ for all $i \in [0, 1]^T$

The first part of the definition represents the set of task-specific productivity parameters for a

given state of AI technology, T . The second part represents, for each worker $i \in [0, 1]$, the positive measure subset of tasks on the task domain $[0, 1]$ for which the worker meets the skills requirements for a given state of AI technology, T .

We assume that workers—who are taken to be solely motivated by monetary incentives—are free to move to the task that pays them the highest wage (i.e., there are no costs to mobility). Although workers make discrete choices over tasks, given the continuum of workers, we assume that these choices aggregate to a measurable skill-supply correspondence across tasks. As in Model 1, we use the notation $s(\tau)$ for the density of skill in task τ for all $\tau \in [0, 1]$. To ensure that aggregate output does not collapse to zero, we again make the technical assumptions that $\theta(\tau) > 0$ almost everywhere, $\int_0^1 \ln \theta(\tau) d\tau > -\infty$, $s(\tau) > 0$ almost everywhere and $\int_0^1 \ln s(\tau) d\tau > -\infty$.²¹

An equilibrium is a schedule of task prices $\left([p : \tau \mapsto p(\tau) \in (0, \infty)]^T\right)$, a schedule of task outputs $\left([z : \tau \mapsto z(\tau) \in (0, \infty)]^T\right)$, a schedule of task wages $\left([w : \tau \mapsto w(\tau) \in (0, \infty)]^T\right)$, a schedule of skill/worker densities $\left([s : \tau \mapsto s(\tau) \in (0, \infty)]^T\right)$ satisfying $\int_0^1 s(\tau) d\tau = 1$, and an associated aggregate productivity, $\bar{A} > 0$, such that all (task output and skill) markets clear and no worker has an incentive to move from their current task.

An equilibrium exists under standard regularity conditions. Consider the set of admissible task-wage schedules, defined as bounded, strictly positive, measurable functions over the continuum of tasks. Given a task-wage schedule, workers choose among the tasks for which they meet the skills requirements in order to maximize wages (i.e., each worker, i , chooses a highest-wage option from the set $\Omega(i)$). Although individual workers may be indifferent across multiple tasks, the atomless nature of the workforce implies that these discrete choices aggregate into a well-behaved skill-supply schedule across tasks: arbitrarily small fractions of workers can be allocated among equally attractive tasks without affecting individual optimality.

On the production side, firms' demand for skills in each task is well defined given wages and the task-based production technology. Market clearing requires that aggregate skill supply equal aggregate skill demand almost everywhere. Taken together, worker optimality and market clearing define a correspondence that maps admissible wage schedules into themselves. This correspondence has nonempty, convex values and is continuous in the appropriate sense, ensuring the existence

²¹Following the same basic structure as used in Model 1, we also impose the requirement that the task-specific productivity parameters are bounded from above: i.e., $\theta(\tau) \leq \bar{\theta} < \infty$, for all $\tau \in [0, 1]$.

of a fixed point. At such a fixed point, wages, prices, task allocations, and output are mutually consistent, and a competitive equilibrium exists.²²

The assumed functional form of the task-based production function is identical to Model 1 (see equation (1)). Using the same procedure as used for Model 1, the equilibrium aggregate productivity can be written as:

$$\dot{A} = A \left[\exp \left(\int_0^1 \ln \theta(\tau) d\tau \right) \right] \left[\exp \left(\int_0^1 \ln s(\tau) d\tau \right) \right]; \quad (29)$$

and the task-specific equilibrium wages are:

$$w(\tau) = \frac{\dot{A}}{s(\tau)} \quad \text{for all } \tau \in [0, 1]. \quad (30)$$

Critically, the $s(\tau)$ values are the equilibrium skill densities given the free mobility decisions of all workers given $\Omega(i)$ for all $i \in [0, 1]$. Note that changes in task-specific productivity parameters (i.e. the θ s) only affect wages through their effect on aggregate productivity given the Cobb-Douglas assumption. It follows that relative wages are invariant to changes in these parameters conditional on the equilibrium skill distribution across tasks. However, it is also clear from equation (30) that an AI-induced reallocation of workers across tasks will affect relative wages and thus the inequality of wages across workers.

We next consider how the degree of uniformity in skill densities in the general equilibrium of the model affects both aggregate productivity and wage inequality across workers. Each effect depends on an intuitive—but different—measure of the uniformity of skill densities across tasks. We then look informally at various cases of how an improvement in the AI technology affects specific workers.

Aggregate Productivity

Proposition 2.1. *A redistribution of skills that raises the geometric mean of task-level skill densities raises aggregate productivity.*

The second term in square brackets on the right-hand-side of equation (29) has a convenient

²²The existence argument relies on the standard convexification logic for economies with a continuum of agents and tasks. Individual workers' best-response sets may be non-singleton, but atomlessness ensures that the induced aggregate skill-supply correspondence is convex-valued and upper hemicontinuous. The equilibrium existence result follows from an application of Kakutani's fixed-point theorem to the correspondence mapping admissible task-wage schedules into themselves. Closely related existence arguments for economies with infinitely many commodities and an atomless set of agents are developed in Aumann (1964), Bewley (1972), and Mas-Colell (1975).

interpretation as the Geometric Mean Uniformity Index for skill densities. We denote this index as $\exp(g)$, where:

$$g = \int_0^1 \ln s(\tau) d\tau. \quad (31)$$

With the task domain on $[0, 1]$, the index has a range $(0, 1]$, with 1 indicating perfect uniformity of skill densities across tasks. With this measure of uniformity, and for any measure of task-specific productivity parameters, a more uniform spread of skill densities (as measured by the index in equilibrium) is associated with higher aggregate productivity. The proof is then immediate from (31).

We use Figure 2 to develop intuition for how this uniformity index (and thus productivity) varies with the underlying equilibrium skill density function. Given that the Beta (α, β) density function $(\alpha > 0, \beta > 0)$ is defined on $[0, 1]$, it provides a convenient way to see how increasing uniformity of the skill densities is associated with rising productivity, where Beta $(1, 1)$ is the uniform distribution on $[0, 1]$. The relationship between $\exp(g)$ and the parameters of the Beta density function is:

$$\exp(g) = \frac{\exp(-(\alpha + \beta - 2))}{B(\alpha, \beta)} \quad \text{where } B(\alpha, \beta) = \int_0^1 \tau^{\alpha-1} (1 - \tau)^{\beta-1} d\tau. \quad (32)$$

To vary the density function, we fix β at 1 and calculate the value of $\exp(g)$ for different values of α . In this case, using well-known results for the Beta distribution, we obtain a particularly simple functional form for equation (33): $\exp(g) = \alpha \exp -(\alpha - 1)$. The more that α diverges from 1 the more the density function will diverge from the uniform density. The figure shows how this uniformity index—and thus aggregate productivity—falls the more that α deviates from 1.

We interpret this as a type of *bottleneck effect* on productivity. For any given set of task-specific productivities, aggregate productivity is at its maximum value of unity when skill densities are uniform across tasks. Of course, the actual task-specific skill densities are determined in equilibrium by the task-location decisions of workers given their sets of task opportunities (which are themselves dependent on their human capital). Productivity declines the more we deviate from uniform skill densities as measured by the geometric mean uniformity index. Indeed, if the improvement in AI tools were to lead to sufficiently strong bottlenecks related to the skill densities, it is even possible that aggregate productivity could fall despite all the changes to the task-specific productivity

parameters being non-negative.

Wage Inequality

Proposition 2.2. *A reallocation of skills that lowers the Shannon Differential Entropy of skill densities will raise wage inequality as measured by the Arithmetic Mean-Geometric Mean ratio.*

While the Geometric Mean Uniformity Index allows us to see how skill density reallocations across tasks will affect productivity, it is also useful to see how such reallocations affect wage inequality across workers. As in Model 1, we use the ratio of the arithmetic mean of wages ($\bar{w}$) to the geometric mean of wages (w^g) as our measure of inequality. Given that we have normalized the size of the workforce to 1, $\bar{w}$ is simply equal to aggregate productivity. Our measure of wage inequality is then:

$$\frac{\bar{w}}{w^g} = \frac{\dot{A}}{\exp\left(\int_0^1 s(\tau) \ln w(\tau) d\tau\right)} \quad (33)$$

$$= \frac{\dot{A}}{\exp\left(\int_0^1 s(\tau) \ln\left(\frac{\dot{A}}{s(\tau)}\right) d\tau\right)} \quad (34)$$

$$= \frac{1}{\exp\left(-\int_0^1 s(\tau) \ln s(\tau) d\tau\right)} \quad (35)$$

$$= \exp(-h) \quad (36)$$

where,

$$h = -\int_0^1 s(\tau) \ln s(\tau) d\tau. \quad (37)$$

h will be recognized as Shannon Differential Entropy, although applied here to skill densities rather than probabilities. This proves Proposition 2.2. Importantly, just like aggregate productivity, we can see that the wage inequality measure also depends on the uniformity of the density function although the nature of that dependence differs.²³

We again use Figure 2 to illustrate how wage inequality changes with changes to the Beta distribution where we fix the β parameter at 1 and vary the α parameter. Although, in general,

²³Ignoring the exponential, h can usefully be viewed as an alternative measure of the uniformity of densities to the measure produced by g . They differ, however, in the “weights” placed on log densities in the calculation of the uniformity measure. g places a weight of 1 on each share; whereas, h weights each log density by its density. As discussed further below, in considering the effects changes in the underlying density function, the two measures will differ in how they respond to the relocation of mass along different parts of the task domain.

a closed form solution relating the α and β parameters to $\exp(-h)$ does not exist, we again draw on well-known results for the Beta distribution when the β is fixed at 1, to write the relationship between our wage inequality measure and the α parameter as: $\exp(-h) = \alpha \exp\left(\frac{(1-\alpha)}{\alpha}\right)$. While the Geometric Mean Uniformity index, $\exp(g)$, reaches its *maximum* value when $\alpha = 1$ (i.e., a uniform distribution), *wage inequality*, $\exp(-h)$, reaches its *minimum* value when $\alpha = 1$. While this illustration may make it seem that aggregate productivity and wage inequality will always move inversely, we next show that the relationship between these two variables is not monotonic.

Proposition 2.3. *While aggregate productivity and wage inequality will typically move inversely due to a reallocation of skill densities, the relationship is not monotonic.*

Assume, contrary to the proposition, that there is always an inverse relationship between productivity and wage inequality. The proof of the proposition consists in showing a counterexample. We start by fixing both A and the first term in square brackets in equation (30) at unity to isolate the effects of any reallocation of skill densities across tasks (i.e., effectively ignoring any direct aggregate productivity effects from changes in A or the θ terms). Both productivity and wage inequality are then fully determined by the density function. Typically, productivity ($\exp(g)$) and our measure of wage inequality ($\exp(-h)$) will move inversely as in our Beta distribution illustration, so that a change in the density function that raises productivity will also tend to reduce wage inequality. However, it is possible that for certain changes in the density function that productivity and wage inequality will move in the *same* direction given that they depend on different measures of the uniformity of that function. If we can find such an example, then the relationship between changes in productivity and wage inequality due to AI-induced skill density reallocations is not monotonic.

To help fix intuition, we first illustrate the possibility of such non-monotonicity with a simple discrete analog with five tasks and five share distributions. Table 1 shows the assumed share distributions. For this discrete example, we use the corresponding discrete case formulas for the geometric mean of skill shares, $\exp(G)$, and Shannon Entropy for skill shares, $\exp(-H)$.²⁴ Figure 3 shows the relationship between $\exp(G)$ and $\exp(-H)$ as we move across these distributions. Typically, a rise in $\exp(-H)$ is associated with a fall in $\exp(G)$. However, the move from Distribution 4 to Distribu-

²⁴The discrete-case geometric mean is given by $\exp(G) = \exp\left(\frac{1}{n} \sum_1^n \ln s_\tau\right)$. The discrete-case Shannon Entropy, H , is given by $H = -\sum_1^n s_\tau \ln s_\tau$. The measure of wage inequality is then: $\exp(-H) = \exp\left(\sum_1^n s_\tau \ln s_\tau\right)$. n is set at 5 in our illustrative example.

tion 3 is an exception in this discrete example; for this change in distribution, $\exp(G)$ and $\exp(-H)$ both rise. Notice that the shift in distribution as we move from 4 to 3 involves the peak becoming more concentrated (which drives the rise in wage inequality as measured by $\exp(-H)$), but the severe “near-hole” that we see in Distribution 4 is substantially filled, mitigating the bottleneck and leading to a rise in aggregate productivity as measured by $\exp(G)$.

Table 1: Illustrative Discrete-Tasks Example with Different Distributions of Skill Shares across Tasks

	Task 1	Task 2	Task 3	Task 4	Task 5
Distribution 1	0.20	0.20	0.20	0.20	0.20
Distribution 2	0.10	0.20	0.40	0.20	0.10
Distribution 3	0.10	0.10	0.60	0.10	0.10
Distribution 4	0.01	0.30	0.29	0.20	0.20
Distribution 5	0.04	0.10	0.72	0.10	0.04

We now turn back to the continuous case. It is known that the Geometric Mean Uniformity Index—and thus aggregate productivity—is particularly sensitive to changes in mass where densities are already low. In terms of our model, this reflects sensitivity to bottlenecks in tasks with low skill densities.²⁵ Because the skill densities weight the log of the density in the calculation of Shannon Differential Entropy, the measure is less sensitive to low densities. For aggregate productivity, then, the main risk can be termed “coverage risk,” with significant “near-holes” in the skill coverage along the task distribution leading to the bottlenecking of productivity. On the other hand, as we saw in our discrete example, Shannon Differential Entropy tends to be more sensitive to changes in mass around the peak of the distribution, which leads to sensitivity of our wage inequality measure to the wages of the lowest paid workers given the inverse relationship between skill supplies and wages. The risk here can be termed “concentration risk,” with high concentrations of workers in a relatively small subset of tasks associated with severe wage inequality.²⁶ Again, while aggregate productivity and wage inequality will tend to move inversely for reasonable changes to the skill density function, there is no necessary relationship between how these variables respond to an AI-induced reshuffling of skill densities.²⁷

²⁵We use the term *bottleneck* to refer to tasks with very low equilibrium skill coverage, which disproportionately depress aggregate productivity under a Cobb–Douglas aggregator.

²⁶We use the term *concentration* to refer to a large mass of workers allocated to a narrow measure of tasks, which raises wage inequality through its effect on the wage distribution.

²⁷Yet another way to see the different effects of shift in the density function is in terms of Kullback-Leibler (KL)

To provide an example of how aggregate productivity and wage inequality can move in the same direction in a continuous task economy, we assume a simple piecewise-constant equilibrium skill density on $[0, 1]$ (see Figure 4). More specifically, consider an economy with a continuum of tasks indexed on the unit interval and an equilibrium distribution of skill supply described by a density function. Partition the task space into three regions: a low-density “dip” region with sparse skill coverage, a center region with moderate density, and a high-density “bump” region in which skills are concentrated. Suppose an improvement in AI tools induces a reallocation of workers out of the center region and into both tails of the distribution. One part of the displaced mass flows into the dip region, improving coverage of previously thinly staffed tasks; another part flows into the bump region, increasing concentration in a narrow set of tasks.

Aggregate productivity depends on the unweighted integral of the log skill density and is therefore especially sensitive to coverage: small increases in skill supply in low-density tasks sharply relax bottlenecks and raise productivity. Wage inequality, by contrast, depends on Shannon Differential Entropy, which weights the log density by the density itself and is therefore most sensitive to concentration in high-density tasks. By choosing the reallocation so that the productivity gains from improved coverage dominate the productivity losses from increased concentration, while the entropy losses from increased concentration dominate the entropy gains from improved coverage, the AI-induced redistribution leads to simultaneous increases in aggregate productivity and wage inequality.

This example establishes non-monotonicity by counterexample. Although productivity and wage inequality will often move inversely following AI-induced reallocations of skills across tasks, the relationship is not monotonic: the same equilibrium redistribution can raise both outcomes because productivity is constrained by coverage across tasks, while inequality is driven by concentration within tasks.

The foregoing example was based on the stringent assumption that the improvement in the AI technology did not change the task-specific productivity parameters—all effects came through the changes in the equilibrium skill densities that resulted from changes to skill requirements. The

divergence. The Geometric Mean Uniformity index is a measure of the divergence of the actual density from the uniform density. The precise form of the relationship between the Geometric Mean Uniformity index and the measure of KL divergence from the uniform is: $\ln(\text{Geometric Mean Uniformity Index}) = -KL(u \parallel s)$. In contrast, Shannon Differential Entropy is a measure of the deviation of the uniform density from the actual density. The relationship is: $\text{Shannon Differential Entropy} = -KL[s \parallel u]$. As is well known, KL divergence is not symmetric.

possibility of aggregate productivity and wage inequality both increasing with an improvement in the AI technology increases once we allow for changes to the task-specific productivity parameters. Conditional on the skill-allocation schedule $s(\tau)$, wage inequality is invariant to the $\theta(\tau)$ schedule. Therefore the $\theta(\tau)$ schedule matters only insofar as it changes the equilibrium $s(\tau)$ schedule. But such changes can affect aggregate productivity “directly” as well as “indirectly” through the equilibrium skill densities (see equation (30)). Such direct effects increase the chances that an improvement in the AI technology will increase aggregate productivity (combining the direct and indirect effects) and raise wage inequality (which is affected only through the indirect effect on the equilibrium schedule of skill densities). Such a situation may arise where the improved AI leads to a greater concentration of equilibrium skill densities as it strongly complements workers with certain specific skills. Productivity is highly sensitive to thinly staffed regions of the task space: reallocations that reduce skill supply in a small subset of tasks can have a disproportionately large negative effect on aggregate productivity through the log aggregator, even though admissibility rules out outright task abandonment. However, when we also allow for the *direct* effect of increased task-specific productivity parameters, productivity is more likely to rise overall, making it more likely that we will observe the joint occurrence of higher aggregate productivity and higher wage inequality.

We next explicitly contrast the cases of “task-opening” developments in AI technology and “task-closing” developments in terms of their likely effects on aggregate productivity and wage inequality.

Task Opening: Consider an improvement in AI tools that reduces the specificity of skill requirements across a broad range of tasks—for example, decision-support systems that allow workers with general analytical and communication skills to perform tasks previously requiring narrow technical expertise. Formally, this corresponds to an expansion of the feasible task sets $\Omega(i)$ for many workers, without disproportionately favoring any particular region of the task space. In equilibrium, such a change induces a reallocation of workers toward previously thinly staffed tasks, flattening the equilibrium skill-density schedule $s(\tau)$.

In this case, the geometric mean uniformity of skill densities increases, raising aggregate productivity by alleviating near-bottlenecks, while Shannon differential entropy also increases, reducing wage inequality. Productivity and inequality move inversely, consistent with a “broadening” effect of AI that rewards general skills.

Task Closing: In contrast, consider an AI improvement that is strongly complementary to a

narrow set of specialized skills—such as AI systems that dramatically increase the productivity of scientists capable of formulating high-quality hypotheses or engineers able to specify complex design constraints. In this case, the availability of AI tools effectively raises the skill threshold for participating in certain tasks, shrinking the feasible task sets $\Omega(i)$ for many workers.

The resulting equilibrium reallocation concentrates skill supply into a relatively small subset of tasks. This concentration lowers Shannon differential entropy, raising wage inequality, and may even reduce geometric mean uniformity by creating thinly staffed regions elsewhere in the task space, potentially lowering productivity through bottleneck effects.

Crucially, under the Cobb–Douglas aggregator, changes in task-specific productivity parameters affect relative wages only through their impact on the equilibrium allocation of skills across tasks. As a result, when task-closing AI also raises task-specific productivity parameters in the concentrated region, aggregate productivity is more likely to increase even as wage inequality increases. This two-channel structure substantially enlarges the set of circumstances under which productivity and inequality move in the same direction relative to the case in which technological change operates solely through reallocation.

Summing up, Model 2 allows for a rich menu of possibilities for how aggregate productivity and wage inequality will respond following an improvement in the AI technology. Critically, the actual effects on these aggregate outcomes will depend on the precise form of the technology shock—e.g., whether it is generally more task-opening or task-closing—and how this shock is intermediated by the human capital stocks of workers. For any given worker, i , their human capital will determine how the set $\Omega(i)$ —i.e., the set of tasks for which they meet the skill requirements—changes with the improvement in the AI technology. We discuss the possible implications for human capital policy in Section 4. Before doing so, we briefly consider the implications of access to improved AI tools for individual workers.

Impacts of an Improvement in AI technology on specific workers: possible cases

While our model allows us to predict the effects of an improvement in AI technology on aggregate productivity and wage inequality, it is also useful to consider how *individual* workers are affected. Table 2 shows the possible cases across two dimensions: whether the worker stays in their current task or moves; and whether the worker experiences an increasing or non-increasing wage. The four

cases exhaust the possible effects on individual workers.

Table 2: Possible Individual Mobility and Wage Effects Following a Change in AI Technology

	Wage $\uparrow$	Wage $\downarrow$ or $\leftrightarrow$
Move	Case I	Case IV
Stay	Case III	Case II

Case I captures the case of a worker who moves because of the improvement in AI technology and obtains a higher wage in their new task than in the old one. We can think of the availability of better AI tools as removing entry barriers to particular tasks as well as also possibly improving productivity in those tasks. An example might be a registered nurse being able to engage in certain diagnostic tasks with the help of AI that were previously in the sole domain of doctors, with the result that the wage of the registered nurse increases.

Case II—staying with the current task but suffering a decrease in their wage (or a stagnant wage)—is broadly the mirror image of Case I. Because of reduced entry barriers and productivity enhancements, the worker (say a doctor engaged in diagnostic tasks) must compete with more workers capable of doing the task, with the increased competition driving down the wage despite any productivity improvements. Even with the wage reduction, however, their current task offers the highest wage from all their task options so that they do not have an incentive to move.

Case III—staying in their current task but benefiting from an increase in their wage—reflects the case of a worker who has their productivity enhanced because of better AI tools and this effect dominates any change in the skill supply for the task. There are two possibilities for the new supply of skills to the task. First, there may be an increased supply due to reduced entry barriers, but the aggregate productivity effect outweighs the supply effect and the wage rises. Second, there is actually a reduced supply of skills so that both the aggregate productivity effect and supply effect both work to increase the wage. An example of the latter might be situation where the availability of AI tools has a complementary effect on highly specialized expertise. For example, we might think of a biomedical scientist who has access to new tools for predicting the binding efficacy of small molecules with some target protein. These tools might allow a given scientist to consider a much larger number of hypotheses for small molecule-protein binding, actually increasing the importance of the scientist’s skill in screening those molecules based on their knowledge of biology that the AI

does not capture.

Case IV—moving to a new task but suffering a decrease in their wage (or a stagnant wage)—reflects the case of a worker who is crowded out of their current task because of an inflow of new competitors, with the wage in the current task driven down despite any increase in aggregate productivity. However, this worker can receive a higher wage in a different task than the new wage in their current task and so moves. Even with this move, they suffer a wage decline (or stagnant wage) compared to the wage they were receiving before the availability of the new AI tools. In contrast to Case III, an example might be an experimental scientist, whose specialized expertise becomes less important because of the access to AI tools, with the resulting increased competition driving them from their current task into a lower paying task.

4 Discussion

We have argued that human workers are likely to retain distinctive capabilities well into the Age of AI. Furthermore, we tentatively located these enduring capabilities in human capacities to make judgments of causal, normative and creative significance. Our assumption of the continued importance of workers in the economy is *not* due to an expectation that the rate of improvement in AI technology will slow—indeed this rate of improvement may remain spectacular. Instead, it is based on what we see as the remarkable capabilities of human workers to use their integrated world models to make judgments of what is significant in the world. In our models, workers thus remain essential across tasks even as the skill requirements needed for tasks change as new AI tools become available. Even though there is no replacement of workers by machines in our models, we have seen that development in AI can still have transformative effects on both aggregate productivity and wage inequality—effects that are intermediated by workers’ human capital stocks. Although our task-based models are highly stylized, we next briefly discuss their possible implications for human capital development policies in a world of ever-improving AI tools.

Recent advances in the economics of human capital have substantially enriched our understanding of the varied forms of human capital that are necessary for success in modern labor markets (see Deming and Silliman, 2025, for an excellent recent review). One result of these advances is an increased emphasis on what Deming calls “higher-order skills” (see also Deming, 2022). The

idea of higher-order skills draws on Bloom’s classic pyramid, which in its modern form sets out an ordering from the base to the top of: remember, understand, apply, analyze, evaluate and create. We think of the importance of human judgment as generally increasing as we ascend the pyramid.²⁸ Decision-making skills—particularly when associated with leading teams—have received increasing emphasis, reflecting the increased importance of such skills in increasingly information-rich economy. A major open question is how human capital policies should respond to a world with ever more capable AI tools. Although our task-based models are written at a high level of abstraction, they do point to the central challenges: how to equip workers with the expertise to use AI; and how to identify the non-AI skills workers will need to have as AI both opens up and closes off task markets to particular workers.

Model 1 highlights the value of equipping a larger fraction of the workforce with AI expertise where that expertise allows workers to use AI in an expanding set of tasks. In the presence of AI advancement at the extensive margin—i.e., AI becoming applicable to an increasing range of tasks—we have seen that a larger share of AI-expert workers will enhance the productivity gain from AI’s increased scope, while also helping to limit adverse effects on wage inequality where not everyone is able to use AI tools.

Of course, a major limiting factor of the model is the all-or-nothing nature of AI expertise. In reality, there are numerous forms of AI expertise—ranging, say, from being able to prompt an LLM to being able to design complex APIs, to being able to fine-tune existing foundation models for enterprise-specific tasks. To capture such richness, the model would need to be extended to allow for the breadth and depth of AI skills that are relevant so as to identify the most advantageous investments in AI expertise. Even so, treating AI as a tool that augments workers draws attention to a set of policy choices that might receive inadequate attention under a pure automation view.

One important way that a worker can build AI expertise is through learning by experimenting.

²⁸A standard criticism of a Bloom-type approach is that it assumes a linear structure in cognitive tasks, with cognitive tasks at higher levels of the pyramid building on tasks at lower levels. As set out in Section 2, we view a worker as using an *integrated* world model to make judgments about the causal and normative significance of events and actions in the world, which suggests a more dynamic interplay between the levels in the Bloom pyramid. Even so, the type of higher-order tasks captured at the upper levels of the pyramid are likely to be especially important in determining the judgment capabilities of the worker. A Bloom-type framework also focuses on the nature of cognitive tasks rather than on a worker’s cognitive capabilities per se. A framework such Howard Gardner’s multiple intelligences—linguistic, logical-mathematical, spatial, bodily-kinesthetic, musical, interpersonal, intrapersonal—focuses directly on these capabilities (Gardner, 1983). Combining these frameworks may provide a more integrated picture of cognitive tasks and the cognitive capabilities that are used to complete them. We can thus talk about the higher-order *skills* necessary to complete higher-order tasks.

Mollick (2024) has emphasized the “jagged” nature of the AI frontier. Even what seem like closely related tasks may offer quite different scopes for effectively using AI. Mollick’s advice is to try using AI for everything the worker does. This allows them to learn about AI by doing and the opportunity to discover where particular forms of AI are applicable. For example, users of AI are experimenting with new work designs to improve the productivity gains from integrating AI into their workflows (e.g., experimenting with how best to integrate AI into work teams: Caplin et al., 2024; Dell’Acqua et al., 2023; Weidmann et al., 2025).

Model 2 highlights a quite different set of human-capital investment decisions. As we have seen, among the important consequences of developments in AI tools are how they open up or close off tasks (or occupations) to workers with particular skills. In the model, to the extent that AI tools on balance open up tasks to workers who previously did not meet the skill requirements for those tasks, this leads to a more uniform density of skill supplies across tasks. Although we have seen that the precise measure of uniformity matters, the effect of opening up occupations will typically be to raise aggregate productivity and reduce wage inequality. This suggests the value of investing in general skills that—with access to AI—are applicable to multiple tasks. Such skills might include communication, critical thinking, teamwork and creative problem-solving capabilities that are widely applicable across tasks. However, we have also seen that access to AI tools might also close off tasks where there is strong complementarity between task-specific skills and the relevant AI tools. In science, for example, the scientific judgment of the scientist—which is likely to depend on the size and integrated nature of their knowledge stock—may be strongly leveraged by having access to AI tools. There would then be a high return to investing in the scarce skills that make particular tasks bottlenecks or sources of wage premiums. The optimal policy might then be a mix of general skill training along with targeted investments in the bottlenecked tasks.²⁹

We have made the simplification of treating AI and non-AI skill acquisition as fully independent. A major issue for the education sector is that the use of AI in the learning process—which is presumably beneficial for the development of AI expertise—might actually hamper the development

²⁹There is a danger in drawing the distinction between general and specific skills too sharply. Epstein (2019) argues for the importance of the worker’s *range* of skills, many of which could be quite specific. He contrasts, for example, an *I* model where the worker has deep skill in one area with a *T* model where the worker has a range of skills, with deep skill in one area (or a *II* model where the worker has a range of skills with deep skills in two areas). Having such a range of skills should give a worker more options in relocating to different tasks as AI changes the skill requirements for tasks/jobs. Given a combinatorial view of creativity, having a range of skills/knowledge could be essential to the worker’s creative capability in any given task.

of other skills. The most discussed case is that of students getting an LLM to do their writing for them. To the extent that the hard work of writing is essential to developing critical thinking skills—which may be just the kind of higher-order skill that would benefit over a career from a pairing with AI—early access to certain types of AI in education may short-circuit key parts of the learning process. Educators and students thus face the difficult task of figuring out how to encourage facility with AI technology without having access to AI impede broader skill acquisition.

Another way in which access to AI might actually hamper skill application and skill development is through its impact on our attention. Attention is a scarce resource that is critical to good judgment, including the type of judgments needed to build our model of the world. Indeed, there is a close parallel between knowing what to attend to and being able to judge what is significant in the world. While our tendency to be surprised may sometimes be the result of inattention, the type of *Peircean surprise* that drives creativity depends on being sufficiently attentive to know when the presence or absence of something demands an explanation or suggests a way that an absent thing might be brought about. There is growing concern that digital tools—including AI, or AI-enhanced digital products—sometimes distract us from seeing the truly significant. To ensure that AI truly serves as productivity-enhancing tools, it is essential that we remain engaged and purposeful users of these tools.

5 Conclusion

The task-based model provides a rich framework for studying the effects of technological change on labor markets. In its canonical form, it hypothesizes the replacement of workers by machines, making it an attractive choice for studying the economic effects of AI on cognitive tasks. Moreover, if technology is on a path to AGI, the model predicts transformative economic effects on aggregate productivity and income distribution, although the nature of the predicted effects depends significantly on whether it is the economy’s final-good production function and/or knowledge production function that gets transformed.

While it is of course impossible to be certain about the future of AI’s development, the early evidence suggests AI is primarily being used as a tool to augment workers rather than as an outright replacement for them. We have located this continued relevance of workers in the enduring

importance of human judgment in most tasks. However, even under a pure augmentation view, the task-based approach remains a powerful framework for studying the effects of AI technology on productivity and income distribution. Although workers remain essential in the completion of tasks in our models, we allow the availability of AI tools to change the skills required for those tasks. To explore the implications of the AI-as-a-tool view in tractable settings we abstracted from capital. We instead treat AI technology as a public good that is freely available to users. Importantly, the models show that AI can still have transformative effects on productivity and income distribution even in a world where workers remain fully employed.

Specifically, we develop two models to highlight the key policy issues in the simplest possible settings. The first model focuses on the scarcity of AI expertise, where technological progress occurs at the extensive margin by expanding the range of tasks that can leverage AI tools. In this setting, we learn that the supply of AI-expert workers acts as a critical lever: a larger share of experts not only amplifies the aggregate productivity gains from new AI tools but also attenuates the rise in wage inequality that typically follows technological shocks. The second model, by contrast, assumes AI expertise is universal and instead isolates the role of workers’ heterogeneous non-AI skills. Here, the main effect of AI is to alter the feasibility of tasks for different workers, driving a reallocation of labor that reveals a decoupling of aggregate outcomes: we find that productivity is driven by the alleviation of bottlenecks (geometric mean uniformity), while wage inequality is driven by the concentration of workers in specific tasks (Shannon differential entropy), implying that AI can sometimes increase both simultaneously depending on how it reshapes the landscape of non-AI skill requirements.

From a policy perspective, the AI-as-tool view highlights the importance of human capital; a form of productive capital that would be of declining importance if we were really on a path to a world in which AI was superior to humans in all cognitive tasks. Our two models highlight the key policy issues in the simplest possible settings. Our results suggest a fundamental divergence in optimal policy strategy relative to the replacement-centered view. Whereas the replacement framework logically points toward defensive redistributive measures, such as universal basic income (UBI), to manage the obsolescence of labor, the tool-based framework mandates an offensive investment in human capital. Specifically, our first model implies that maximizing the welfare gains from AI requires policies that broaden the distribution of AI-specific expertise; expanding the share of the

workforce capable of deploying these tools amplifies aggregate productivity gains while simultaneously mitigating the exclusionary wage premiums that would otherwise accrue to scarce experts. Furthermore, our second model highlights that if AI tools induce “task concentration”—thereby exacerbating inequality—policy may beneficially prioritize the development of complementary higher-order skills, such as complex judgment and evaluation, which allow workers to migrate into and alleviate productivity bottlenecks. Consequently, the economic impact of AI will be determined not solely by the technological frontier, but by the elasticity of the skill supply in response to shifting task requirements.

Our view of human capital is influenced by recent developments in the human capital literature. Given the increasing complexity of modern economies, more attention is now given to so-called “higher-order skills,” which include such factors as the capability to make good decisions and the capability to act effectively in teams. We see this as paralleling our emphasis on human judgment in complex physical and social settings. The focus is on how AI can aid this judgment rather than replacing it. How human capital policies should respond to world in which AI tools are becoming increasingly important judgment-support tools remains largely an open question.

Our models are highly stylized by design to focus on what we see as neglected human capital-intermediated mechanisms. Important directions for future work include the integration of the AI-as-tool and AI-as-replacement approaches, which will require explicit consideration of capital in the model. We have also assumed a CRS Cobb-Douglas aggregator for tractability. It will be important to explore the implications of richer substitution possibilities. Finally, although separating out the AI-expertise and non-AI skills dimensions in separate (though related) models allowed us to derive key results in relatively transparent settings, future work will put them together in a reasonably tractable integrated model.

References

- Acemoglu, D. (2025). The simple macroeconomics of AI. *Economic Policy* 40(121), 13–58.
- Acemoglu, D. and D. Autor (2011). Skills, tasks and technologies: Implications for employment and earnings. In O. Ashenfelter and D. Card (Eds.), *Handbook of Labor Economics*, Volume 4, pp. 1043–1171. Amsterdam: Elsevier.
- Acemoglu, D. and S. Johnson (2023). *Power and Progress: Our Thousand-Year Struggle Over Technology and Prosperity*. London: Basic Books.
- Acemoglu, D. and P. Restrepo (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review* 108(6), 1488–1542.
- Acemoglu, D. and P. Restrepo (2019). Artificial intelligence, automation, and work. In A. Agrawal, J. Gans, and A. Goldfarb (Eds.), *The Economics of Artificial Intelligence: An Agenda*, pp. 197–236. Chicago: University of Chicago Press.
- Aghion, P., B. F. Jones, and C. I. Jones (2019). Artificial intelligence and economic growth. In A. Agrawal, J. Gans, and A. Goldfarb (Eds.), *The Economics of Artificial Intelligence: An Agenda*, pp. 237–282. Chicago: University of Chicago Press.
- Agrawal, A., J. Gans, and A. Goldfarb (2018a). Prediction, judgment and complexity: A theory of decision making and artificial intelligence. In A. Agrawal, J. Gans, and A. Goldfarb (Eds.), *The Economics of Artificial Intelligence: An Agenda*, pp. 89–110. Chicago: University of Chicago Press.
- Agrawal, A., J. Gans, and A. Goldfarb (2018b). *Prediction Machines: The Simple Economics of Artificial Intelligence*. Boston: Harvard Business Review Press.
- Agrawal, A., J. Gans, and A. Goldfarb (2019a). Artificial intelligence: The ambiguous labor market impact of automating prediction. *Journal of Economic Perspectives* 33(2), 31–50.
- Agrawal, A., J. Gans, and A. Goldfarb (2019b). Exploring the impact of artificial intelligence: Prediction versus judgment. *Information Economics and Policy* 47, 1–6.
- Agrawal, A., J. Gans, and A. Goldfarb (2022). *Power and Prediction: The Disruptive Economics of Artificial Intelligence*. Boston: Harvard Business Review Press.
- Agrawal, A., J. Gans, and A. Goldfarb (2025). The economics of bicycles for the mind. NBER Working Paper 34034, National Bureau of Economic Research.
- Agrawal, A., J. McHale, and A. Oettl (2019). Finding needles in haystacks: Artificial intelligence and recombinant growth. In A. Agrawal, J. Gans, and A. Goldfarb (Eds.), *The Economics of Artificial Intelligence: An Agenda*, pp. 149–174. Chicago: University of Chicago Press.
- Agrawal, A., J. McHale, and A. Oettl (2023). Superhuman science: How artificial intelligence may impact innovation. *Journal of Evolutionary Economics* 33, 1473–1517.
- Agrawal, A., J. McHale, and A. Oettl (2024). Artificial intelligence and scientific discovery: A model of prioritized search. *Research Policy* 53(5), 104989.
- Agrawal, A., J. McHale, and A. Oettl (2026). AI and science. In M. Macgarvie and R. Veugelers (Eds.), *The Economics of Science*. Chicago: University of Chicago Press. Forthcoming.

- Aumann, R. J. (1964). Markets with a continuum of traders. *Econometrica* 32(1–2), 39–50.
- Autor, D. (2024, February). AI could actually help rebuild the middle class. *Noēma*. Berggruen Institute.
- Autor, D. and N. Thompson (2025). Expertise. *Journal of the European Economic Association* 23(4), 1203–1271.
- Autor, D. H., F. Levy, and R. J. Murnane (2003). The skill content of recent technological change: An empirical exploration. *Quarterly Journal of Economics* 118(4), 1279–1333.
- Bewley, T. F. (1972). Existence of equilibria in economies with infinitely many commodities. *Journal of Economic Theory* 4(3), 514–540.
- Brynjolfsson, E. (2022). The turing trap: The promise & peril of human-like artificial intelligence. *Dædalus* 151(2), 272—287.
- Brynjolfsson, E., D. Li, and L. R. Raymond (2025). Generative AI at work. *Quarterly Journal of Economics* 140(2), 889–942.
- Caplin, A., D. J. Deming, S. Li, D. J. Martin, P. Marx, B. Weidmann, and K. J. Ye (2024, October). The ABC’s of who benefits from working with AI: Ability, beliefs, and calibration. NBER Working Paper 33021, National Bureau of Economic Research.
- Chatterji, A., T. Cunningham, D. J. Deming, Z. Hitzig, C. Ong, C. Y. Shan, and K. Wadman (2025). How people use ChatGPT. NBER Working Paper 34255, National Bureau of Economic Research.
- Chollet, F. (2019). On the measure of intelligence. arXiv:1911.01547.
- Cunha, F. and J. J. Heckman (2007). The technology of skill formation. *American Economic Review* 97(2), 31–47.
- Dell’Acqua, F., E. McFowland, E. Mollick, H. Lifshitz-Assaf, K. C. Kellogg, S. Rajendran, L. Kraymer, F. Candelon, and K. R. Lakhani (2023, September). Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality. Working Paper 24-013, Harvard Business School.
- Deming, D. J. (2017). The growing importance of social skills in the labor market. *Quarterly Journal of Economics* 132(4), 1593–1640.
- Deming, D. J. (2022). Four facts about human capital. *Journal of Economic Perspectives* 36(3), 75–102.
- Deming, D. J. and M. Silliman (2025). Skills and human capital in the labor market. In D. Card and O. Ashenfelter (Eds.), *Handbook of Labor Economics*, Volume 6, pp. 115–152. Amsterdam: Elsevier.
- Epstein, D. (2019). *Range: How Generalists Triumph in a Specialized World*. New York: Riverhead Books.
- Gardner, H. (1983). *Frames of Mind: A Theory of Multiple Intelligences*. New York: Basic Books.


- Harnad, S. (1990). The symbol grounding problem. *Physica D: Nonlinear Phenomena* 42(1–3), 335–346.
- Heckman, J. J. (2013). *Giving Kids a Fair Chance (A Strategy That Works)*. Boston Review Books. Cambridge, MA: MIT Press.
- Heckman, J. J. and C. O. Corbin (2016). Capabilities and skills. *Journal of Human Development and Capabilities* 17(3), 342–359.
- Heckman, J. J. and T. Kautz (2012). Hard evidence on soft skills. *Labour Economics* 19(4), 451–464.
- Jones, B. F. (2025). Artificial intelligence in research and development. NBER Working Paper 34312, National Bureau of Economic Research.
- Kahneman, D. (2011). *Thinking, Fast and Slow*. London: Penguin Books.
- Kahneman, D. (2019). Comment on colin camerer, “artificial intelligence and behavioral economics”. In A. Agrawal, J. Gans, and A. Goldfarb (Eds.), *The Economics of Artificial Intelligence: An Agenda*, pp. 609–612. Chicago: University of Chicago Press.
- Lake, B. M., T. D. Ullman, J. B. Tenenbaum, and S. J. Gershman (2017). Building machines that learn and think like people. *Behavioral and Brain Sciences* 40, e253.
- LeCun, Y. (2022). A path towards autonomous machine intelligence. OpenReview. Version 0.9.2.
- Mas-Colell, A. (1975). A model of equilibrium with differentiated commodities. *Journal of Mathematical Economics* 2(2), 263–295.
- Mollick, E. (2024). *Co-Intelligence: Living and Working with AI*. New York: Portfolio.
- Nagel, T. (1986). *The View from Nowhere*. New York: Oxford University Press.
- Noy, S. and W. Zhang (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science* 381(6654), 187–192.
- Peirce, C. S. (1998). *The Essential Peirce: Selected Philosophical Writings, Volume 2 (1893–1913)*. Bloomington and Indianapolis: Indiana University Press. Edited by the Peirce Edition Project; originally written 1903.
- Peng, S., E. Kalliamvakou, P. Cihon, and M. Demirer (2023). The impact of AI on developer productivity: Evidence from GitHub Copilot. arXiv:2302.06590.
- Polanyi, M. (1966). *The Tacit Dimension*. Chicago: University of Chicago Press.
- Putnam, H. (2002). *The Collapse of the Fact/Value Dichotomy and Other Essays*. Cambridge, MA: Harvard University Press.
- Restrepo, P. (2025). We won’t be missed: Work and growth in the AGI world. NBER Working Paper 34423, National Bureau of Economic Research.
- Searle, J. R. (2004). *Mind: A Brief Introduction*. New York: Oxford University Press.
- Sen, A. (1987). *On Ethics and Economics*. Oxford: Basil Blackwell.

- Silver, D. and R. S. Sutton (2025). Welcome to the era of experience. In G. Konidaris (Ed.), *Designing an Intelligence*. Cambridge, MA: MIT Press. Preprint available at <https://davidstarsilver.wordpress.com/publications/>.
- Weidmann, B., Y. Xu, and D. J. Deming (2025, April). Measuring human leadership skills with AI agents. NBER Working Paper 33662, National Bureau of Economic Research.
- Zeira, J. (1998). Workers, machines, and economic growth. *Quarterly Journal of Economics* 113(4), 1091–1117.


Figures

Figure 1

(a) Potential Gain from an Improvement in AI Technology on the Extensive Margin



(b) Impact of an Increase in the AI-Expert Share of the Workforce on the Gain from an improvement in the AI Technology



Panel (a) shows the constrained comparative advantage equilibrium where the production of new knowledge requires a continuum of tasks indexed from 0 to 1. $\bar{\tau}$ denotes the share of tasks for which AI can be deployed, where $\bar{\tau}$ is viewed as a technological constraint given the current state of AI. The use of AI requires complementary AI-expert knowledge workers. It is assumed that the current equilibrium is constrained in the sense that the marginal task would be reallocated to AI-expert workers if it were technologically feasible. As shown, the gap between the curves is a measure of the productivity gain from being able to use AI-expert scientists for a marginal task (see text). Technological improvements in AI are therefore associated with increased knowledge productivity. Panel (b) shows how the size of the gain from a marginal improvement in the share of tasks for which it is feasible to use AI-expert workers will increase with the share of AI-experts in the workforce. The mechanism is that increase in the relative supply of AI-expert workers drives down their relative wage, thus decreasing the cost and increasing the productivity gains from employing these workers in equilibrium.

Figure 2: $\exp(g)$ and $\exp(-h)$ for Various Beta ($\alpha, \beta = 1$) Distributions




Figure 3: Non-Monotonic Relationship between Wage Inequality ($\exp(-H)$) and Aggregate Productivity ($\exp(G)$) for the Skill-Share Distributions Shown in Table 1





Figure 4: Illustrative Continuous-Tasks Example with a Partitioned Task Range



Notes: The figure shows the equilibrium skill densities before (thin solid lines) and after (thick dashed lines) a change in AI technology. The task range is partitioned into three regions: a low-density dip region; a high-density bump region; and a middle-density center region. The change in technology moves mass from the center region to both the dip and bump regions. The dip region has a disproportionate effect on aggregate productivity given $\exp(g) = \exp\left(\int_0^1 \ln s(\tau) d\tau\right)$. Overall, the reallocation of mass from the center to the dip and bump regions is assumed to *increase* aggregate productivity as uniformity *increases* as measured by Geometric Mean Uniformity. Conversely, the bump region has a disproportionate effect on wage inequality given $\exp[-h] = \exp\left(\int_0^1 s(\tau) \ln(s(\tau)) d\tau\right)$. Overall, the reallocation of mass from the center to the dip and bump regions is assumed to *increase* wage inequality as uniformity *decreases* as measured by Shannon Differential Entropy. Reflecting the different ways that skill densities are weighted under the two measures, the reallocation of densities is thus assumed to increase both aggregate productivity and wage inequality in this example.