

NBER WORKING PAPER SERIES

GENIUS ON DEMAND:
THE VALUE OF TRANSFORMATIVE ARTIFICIAL INTELLIGENCE

Ajay K. Agrawal
Joshua S. Gans
Avi Goldfarb

Working Paper 34316
<http://www.nber.org/papers/w34316>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2025

Thanks to Claude Sonnet 3.7 for helpful research assistance and to Luis Garicano and participants at the NBER Economics of Transformative AI Conference (Stanford, 2025) for useful comments. Responsibility for all errors remains our own. 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/w34316>

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2025 by Ajay K. Agrawal, Joshua S. Gans, and Avi Goldfarb. 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.

Genius on Demand: The Value of Transformative Artificial Intelligence
Ajay K. Agrawal, Joshua S. Gans, and Avi Goldfarb
NBER Working Paper No. 34316
October 2025
JEL No. D24, J24, O33

ABSTRACT

This paper examines how the emergence of transformative AI systems providing "genius on demand" would affect knowledge worker allocation and labour market outcomes. We develop a simple model distinguishing between routine knowledge workers, who can only apply existing knowledge with some uncertainty, and genius workers, who create new knowledge at a cost increasing with distance from a known point. When genius capacity is scarce, we find it should be allocated primarily to questions at domain boundaries rather than at midpoints between known answers. The introduction of AI geniuses fundamentally transforms this allocation. In the short run, human geniuses specialise in questions that are furthest from existing knowledge, where their comparative advantage over AI is greatest. In the long run, routine workers may be completely displaced if AI efficiency approaches human genius efficiency.

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

Avi Goldfarb
University of Toronto
Rotman School of Management
and NBER
agoldfarb@rotman.utoronto.ca

Joshua S. Gans
University of Toronto
Rotman School of Management
and NBER
joshua.gans@rotman.utoronto.ca

1 Introduction

Recent advances in artificial intelligence (AI) have raised the prospect of machines with extraordinary cognitive capabilities—systems that could fundamentally reshape labour markets by providing what Amodei (2024) describes as a “country of geniuses in a data center.” Such systems, capable of applying and potentially creating knowledge across various domains, raise a fundamental economic question: How would labour markets absorb such an influx of cognitive capacity, and what would be the consequences for different types of knowledge workers?

In order to analyse this question, it is important to distinguish geniuses from other knowledge workers. In this regard, we are inspired by this story concerning the approach taken by John von Neumann.

Then there is the famous fly puzzle. Two bicyclists start twenty miles apart and head toward each other, each going at a steady rate of 10 m.p.h. At the same time a fly that travels at a steady 15 m.p.h. starts from the front wheel of the southbound bicycle and flies to the front wheel of the northbound one, then turns around and flies to the front wheel of the southbound one again, and continues in this manner till he is crushed between the two front wheels. Question: what total distance did the fly cover? The slow way to find the answer is to calculate what distance the fly covers on the first, northbound, leg of the trip, then on the second, southbound, leg, then on the third, etc., etc., and, finally, to sum the infinite series so obtained. The quick way is to observe that the bicycles meet exactly one hour after their start, so that the fly had just an hour for his travels; the answer must therefore be 15 miles. When the question was put to von Neumann, he solved it in an instant, and thereby disappointed the questioner: “Oh, you must have heard the trick before!” “What trick?” asked von Neumann; “all I did was sum the infinite series.” Halmos (1973), 386-287.

While normal mathematicians would use tricks to simplify a problem, von Neumann could work from first principles. In effect, he reproduced the knowledge required to answer the question. Our notion of a genius in this paper is similar. We focus on how knowledge is both created and utilised by knowledge workers (including AIs) to impact decision-making. Routine workers rely on existing knowledge and draw inferences from that to answer questions. Geniuses (human or AI) generate new but ephemeral knowledge to answer questions put to them (working from first principles). Thus, routine workers and geniuses differ fundamentally in how they approach problems. The former can easily guess answers that are imperfectly accurate, while the latter can deliver accurate answers at a cost. Moreover, to align with Amodei’s conception, we assume that AI geniuses approach problems in a similar manner to their human counterparts.

Using this distinction, we model how organisations allocate different knowledge worker types to questions that arise. The key idea is that when questions arise, it is important

that someone knows what type of knowledge worker would best address that question. As knowledge workers use two distinctive approaches and also differ in terms of their availability, we view the task of “knowing who to ask” as a key learned task – from experience – and hence, as an equilibrium outcome.¹ In our model, this function is represented by a manager who routes questions to the appropriate type of worker, taking into account both their absolute advantage and, because specifically human geniuses are scarce, their comparative advantage.² We characterise the optimal routing rule that maximises the total value of answers provided. Importantly, our framework shows how routine and genius workers, approaching questions differently, have value generated with distinct curvature implications across the space of possible questions. This drives their allocation to question types.

Following this, using this same setup, we examine the short and long-run impacts of an influx of AI geniuses. Recall that the task of knowing who to allocate a question to is a learned one, and so when AI geniuses become abundant, it takes a while for the manager to reoptimise routing. Thus, the short run is characterised by the manager not altering their pre-AI allocation rule and allocating AI geniuses to genius-type questions along with human geniuses. We show that this pushes human geniuses to tackle more novel questions. In the long run, the AI geniuses are shown to displace routine work as the manager adjusts their routing rule accordingly. The short and long-run framing allows us to consider the path by which AI geniuses might impact knowledge work.

The paper builds on the model of Carnehl and Schneider (2025) that provides a micro-foundation for how knowledge is used in decision-making. Gans (2025b) applies their model to understand the role of AI but as a decision-assistant rather than as a Amodei type genius, and focuses on what this does for the creation of knowledge itself. While here, AI is used as an autonomous genius knowledge worker, it differs from approaches such as Acemoglu and Restrepo (2018), moving beyond a simple substitution of labour for capital. The closest work is that of Ide and Talamas (2025), who similarly consider the impact of AI on knowledge workers. They, however, endow different workers (and AI) with differing degrees of knowledge and explore hierarchical knowledge processing (a la Garicano (2000)) as a foundation, whereas the model here considers agents that process existing knowledge (available to all) in different ways.

The remainder of the paper is organised as follows. Section 2 presents the model, building

¹Garicano (2000) also focusses on knowing who can answer a question in an organisation but his conception relies on who has the requisite information rather than how a worker processes knowledge. Gans (2025a) provides a model closer to the model in this paper.

²In an important sense, the manager is the agent who drives the demand for geniuses by knowing when to optimally route problems to them. One could view this as a demand-side proxy for consumer demand for genius-level problem solving.

on the Carnehl and Schneider (2025) framework of knowledge as question-answer pairs but simplifying to a single knowledge point. Section 3 defines the two types of knowledge workers and their capabilities. Section 4 formulates the manager’s problem and derives the optimal routing rule. Section 5 introduces AI geniuses and analyses their impact on worker allocation in both the short and long run. Section 6 concludes with directions for future research.

2 Model Setup

To analyse how transformative AI might reshape knowledge worker allocation, we require a framework that precisely captures the relationship between knowledge, uncertainty, and value creation. Our approach builds on a microfounded model of knowledge that explicitly connects the structure of what is known to the value of what can be produced. This microfoundation is essential for understanding knowledge work, where value derives not from physical output but from the quality of decisions made under uncertainty.

2.1 Knowledge Structure and Uncertainty

Following Carnehl and Schneider (2025), we model knowledge as a collection of question-answer pairs distributed along the real line. Each question $x \in \mathbb{R}$ has a unique answer $y(x) \in \mathbb{R}$. The true answers to questions are determined by a realisation of a standard Brownian motion, with $y(0) = 0$ as a normalisation. This structure captures the intuitive notion that questions that are “close” to each other tend to have related answers, with the relationship becoming more uncertain as the distance between them increases.

In our simplified model, we assume knowledge consists of a single known question-answer pair, denoted by $\mathcal{F}_t = \{(x_0, y(x_0))\}$. For questions that haven’t yet been answered, agents can form conjectures based on existing knowledge. For a question x located at some distance from the known point x_0 , the conjecture follows a normal distribution with mean:

$$\mu_x(\mathcal{F}_t) = y(x_0) \tag{1}$$

and variance:

$$\sigma_x^2(\mathcal{F}_t) = |x - x_0| \tag{2}$$

This formulation captures a fundamental characteristic of knowledge: uncertainty (variance) is lowest at the known point and increases linearly with distance from the known point. This pattern of uncertainty mirrors real-world scenarios, where extrapolating from established knowledge becomes increasingly difficult as one moves further from what is already known.

2.2 The Value of Knowledge in Decision-Making

The value of knowledge stems from its instrumental use in decision-making. Following the decision-theoretic framework of Carnehl and Schneider (2025), we model the payoff structure for decision-makers as follows: When confronted with a question at position x , a decision-maker must choose whether to take an action $a(x)$ or abstain (denoted as $a(x) = \emptyset$). The payoff from taking an action is defined by a quadratic loss function:

$$u(a(x), x) = \begin{cases} 0, & \text{if } a(x) = \emptyset, \\ 1 - \frac{(a(x) - y(x))^2}{q}, & \text{if } a(x) \in \mathbb{R}. \end{cases} \quad (3)$$

Here $q > 0$ is a critical parameter representing the importance of precision—or alternatively, the strength of the outside option. When the uncertainty associated with a question exceeds the threshold q , decision-makers rationally choose to abstain rather than risk making a highly inaccurate decision.

Given this payoff structure, the optimal decision is to select $a^*(x|\mathcal{F}_t) = \mu_x(y(x)|\mathcal{F}_t)$ (the expected value of the answer) provided that the uncertainty (variance) is sufficiently low: $\sigma_x^2(Y|\mathcal{F}_t) \leq q$. When this condition is met, the expected payoff from applying knowledge is:

$$V(x) = 1 - \frac{\sigma_x^2(y(x)|\mathcal{F}_t)}{q} \quad (4)$$

This value function has important properties: it approaches 1 (maximum value) as uncertainty approaches 0 (perfect knowledge) and decreases linearly toward 0 as uncertainty approaches the threshold q .

2.3 Knowledge Domain and Regional Structure

For analytical tractability, we consider a bounded knowledge domain centred around our single knowledge point x_0 . While the model here can easily accommodate arbitrary domains with x_0 asymmetrically located within them, for notational simplicity, it is assumed that the domain over which questions arrive is uniformly distributed on $[0, 1]$ with $x_0 = \frac{1}{2}$.

We further assume that $\frac{1}{2} > q$, meaning that sufficiently far from the known point (specifically in the regions $[0, \frac{1}{2} - q]$ and $[\frac{1}{2} + q, 1]$), uncertainty exceeds the threshold q and so routine workers would choose not to take an action in these regions.

These assumptions create a natural partition of the knowledge domain into three regions, each with distinct characteristics regarding uncertainty and decision-making:

1. Region 1: $[0, \frac{1}{2} - q]$ - Questions where uncertainty exceeds the threshold q , making

abstention optimal for routine workers.

2. Region 2: $[\frac{1}{2} - q, \frac{1}{2} + q]$ - Questions where uncertainty increases with distance from x_0 but remains below the threshold q .
3. Region 3: $[\frac{1}{2} + q, 1]$ - Questions where uncertainty exceeds the threshold q , making abstention optimal for routine workers.

This regional structure (see Figure 1) creates interesting patterns of comparative advantage between different worker types, as we will explore in the next section.

3 Types of Knowledge Workers

Our economy features three distinct types of economic agents in the knowledge production and application process. First, routine workers possess specialised skills to apply existing knowledge to questions that fall within established frameworks. Second, genius workers have the unique ability to create new knowledge, enabling them to tackle questions that lie beyond the frontier of what is currently understood. Finally, managers serve as allocators who determine which questions should be routed to routine versus genius workers, maximising the overall value of knowledge work subject to capacity constraints. In this section, we elaborate on the capabilities, value functions, and real-world interpretations of each worker type.

3.1 Routine Workers: Application of Existing Knowledge

Routine workers represent intellectual labour that applies existing knowledge to form conjectures about questions. These workers can be thought of as skilled practitioners who operate within established paradigms—they extrapolate from what is known but do not create fundamentally new knowledge. Examples include residents in medicine applying standard differential diagnoses, paralegals applying precedent to routine cases, or software developers implementing known algorithms for familiar problem types.

Mathematically, routine workers follow the decision-making framework described in the previous section. For a question at position x , a routine worker:

1. Forms a conjecture based on existing knowledge, with mean $\mu_x(\mathcal{F}_t)$ and variance $\sigma_x^2(\mathcal{F}_t)$
2. Provides an answer if and only if $\sigma_x^2(\mathcal{F}_t) \leq q$, otherwise abstains
3. When answering, provides the expected value $\mu_x(\mathcal{F}_t)$ as the answer




Figure 1: Decision-Making Value for Routine and Genius Workers

4. Generates a value of $1 - \frac{\sigma_x^2(\mathcal{F}_t)}{q}$ when answering, and 0 when abstaining

Critically, routine workers cannot provide valuable answers to questions in Regions 1 and 3, where uncertainty exceeds the threshold q . Their value contribution decreases linearly with distance from the known point, reaching zero at the boundaries of Region 2.

For a question at position x , the value created by a routine worker is:

$$V_R(x) = \begin{cases} 1 - \frac{|x-x_0|}{q}, & \text{if } |x - x_0| \leq q, \\ 0, & \text{otherwise.} \end{cases} \quad (5)$$

This value is depicted in Figure 1. A key characteristic of this value function is its linear decrease with distance from the known point x_0 . Moreover, we assume that there is an abundance of routine workers, i.e., $L_R \geq 2q$.

3.2 Genius Workers: Creative Knowledge Generation

Genius workers represent intellectual labour that creates new knowledge rather than merely applying existing knowledge. These workers can be thought of as innovators, experts, or creative specialists who generate novel insights rather than extrapolating from what is already known. Examples include specialist physicians diagnosing unusual presentations, senior attorneys developing novel legal theories, or research scientists creating new algorithms for previously unsolved problems.

Unlike routine workers, genius workers can generate exact answers for any question,

including those in Regions 1 and 3 where uncertainty exceeds the threshold for routine workers. However, creating new knowledge incurs a cost that increases with distance from existing knowledge—the further from established knowledge, the more difficult and resource-intensive the creative process becomes.

Genius workers create new knowledge to answer questions. However, their ability to do so is anchored in existing (public) knowledge, x_0 . We model this by assuming that when genius workers receive questions, x , that are increasingly distant from the type that would be answered by existing knowledge, the cost of producing the knowledge to answer the question posed increases. Thus, for a question at position x , the cost of creating knowledge at this point derives from the distance from the known point:

$$\text{Cost}(x) = \eta \cdot (x - x_0)^2 \quad (6)$$

where $\eta > 0$ is a cost parameter that reflects the difficulty of creating new knowledge. This quadratic cost function captures the intuition that creating knowledge becomes progressively more difficult as one moves further from what is already known.³

The net value created by a genius worker answering a question at position x is therefore:

$$V_G(x) = 1 - \text{Cost}(x) = 1 - \eta \cdot (x - x_0)^2 \quad (7)$$

This is also depicted in Figure 1. This quadratic cost structure creates a value function that decreases more rapidly with distance than that of routine workers, but maintains positive contributions across the entire domain for appropriate parameter values. It is assumed here that geniuses can answer questions over the entire domain $[0, 1]$ by assuming that $1 - \eta \frac{1}{4} \geq 0 \Leftrightarrow \eta \leq 4$.⁴

3.3 Marginal Effects from Known Knowledge Distance

The central distinction between routine and genius workers lies not just in their capabilities (application versus creation of knowledge) but critically in the different ways their value

³It is instructive to note that this ‘knowledge creation’ cost function differs from that utilised by Carnehl and Schneider (2025). They consider new knowledge as arising out of a targeted search process and, consequently, the costs of finding knowledge increase with the variance associated with the answer process described earlier. That variance, as we already noted, is concave in the distance from existing knowledge, which is what gives the routine worker value its linear shape. Here, however, we consider the generation of new answers to questions posed as a result of cognitive effort or capacity, and, thus, it is appropriate that the costs of doing so are convex. Nonetheless, the broad conclusions of the paper would be unchanged if genius costs were linear (for human and AI alike).

⁴If η is sufficiently high, then routine workers may strictly dominate geniuses over domains further from x_0 . This would be a more convincing outcome if such workers had access to additional knowledge points as in Gans (2025a).




Figure 2: Absolute advantage and implied allocation. In Region 2, $AA(d) = \frac{d}{q} - \eta d^2$ peaks at $d^* = \min(q, \frac{1}{2\eta q})$; for $d \geq q$, $AA(d) = 1 - \eta d^2$ is maximized at $d = q$. Initial genius bands lie at $x_0 \pm q$ when $\eta \leq 1/(2q^2)$ and at $x_0 \pm d^*$ when $\eta > 1/(2q^2)$; bands widen as L_G increases.

contributions change with distance from the known knowledge point. This difference in value curvature fundamentally shapes the optimal allocation of workers across the knowledge domain:

1. **Routine Worker Value Curvature:** The value function $V_R(x)$ exhibits linear decline with distance from x_0 through the term $\frac{|x-x_0|}{q}$. At the threshold distance q from the known point, their value contribution drops to zero—a precipitous boundary beyond which they cannot contribute.
2. **Genius Worker Value Curvature:** The value function $V_G(x)$ follows a quadratic decline through the term $\eta \cdot (x - x_0)^2$. While their value also decreases with distance from the known point, the decline is initially more gradual than for routine workers near x_0 but more severe at greater distances.

These contrasting value curvatures create a complex pattern of comparative advantage across the knowledge domain; see Figure 2. At points very close to existing knowledge, routine workers may deliver nearly the same value as genius workers, but at presumably lower cost. As distance increases, the relative advantage shifts depending on the specific parameter values, creating distinct regions where each worker type has comparative advantage.

3.4 The Manager: Knowledge Work Allocation

The third agent in our framework is the manager, who plays a central coordinating role in the knowledge economy. The manager can be conceptualised as a manager, dispatcher, or allocation mechanism determining which worker type should handle each incoming question. In modern knowledge organisations, this role might be embodied in triage systems, project management frameworks, or hierarchical structures that route work to appropriate specialists.

Managers possess knowledge of both the question domain (understanding where questions fall relative to existing knowledge) and of worker capabilities (understanding the relative strengths of routine versus genius workers across the domain). This meta-knowledge enables matching of questions to worker types. The manager does not directly answer questions but rather maximises the total value creation of the knowledge system through optimal routing decisions.

In the next section, we will formally define the manager’s optimisation problem and characterise the efficient allocation of knowledge workers across the question domain.

4 Optimal Worker Allocation

In this section, we tackle the central economic problem of our model: how should a manager optimally allocate genius and routine workers across the knowledge domain to maximise total value? This allocation challenge represents a fundamental task in knowledge-intensive organisations—determining which types of intellectual labour should address which questions when both types have constrained capacity.

A manager receives questions drawn uniformly from the knowledge domain $[0, 1]$ and must route each to either a routine worker or a genius worker. The manager has access to L_R routine workers and L_G genius workers, creating capacity constraints on each worker type. We assume that routine workers are in sufficient supply to handle all questions where they can provide positive value (L_R is large enough to cover all of Region 2), but genius workers are scarce ($L_G < 1$).

The manager’s objective is to maximise the total value of answers provided:

$$\max_{A_R, A_G} \int_{A_R} V_R(x) dx + \int_{A_G} V_G(x) dx \quad (8)$$

where A_R and A_G represent the sets of questions assigned to routine and genius workers, respectively. This optimisation is subject to capacity constraints ($|A_R| \leq L_R$ and $|A_G| \leq L_G$) and the requirement that each question can be assigned to at most one worker type

$$(A_R \cap A_G = \emptyset).$$

4.1 Absolute Advantage Across the Knowledge Domain

The key to solving this allocation problem is understanding where genius workers create more value than routine workers. The *absolute advantage* of a genius worker over a routine worker at position x is:

$$AA(x) = V_G(x) - V_R(x) \quad (9)$$

This measures the additional value generated by assigning a genius rather than a routine worker to a question at position x . Significantly, this advantage varies systematically across the knowledge domain, creating a natural ranking of questions by the benefit of genius assignment.

To understand the pattern of absolute advantage, we analyse its behaviour across the three regions of our knowledge domain:

1. **Regions 1 and 3** ($[0, x_0 - q]$ and $[x_0 + q, 1]$): In these regions, routine workers cannot provide answers due to excessive uncertainty ($V_R(x) = 0$). The absolute advantage equals the genius value: $AA(x) = V_G(x) = 1 - \eta(x - x_0)^2$. This advantage decreases with distance from the boundaries $x_0 - q$ and $x_0 + q$.
2. **Region 2** ($[x_0 - q, x_0 + q]$): Here, both worker types can provide answers, but their values follow different curvatures. In Region 2, the absolute advantage is:

$$AA(x) = 1 - \eta(x - x_0)^2 - \left(1 - \frac{|x - x_0|}{q}\right) = \frac{|x - x_0|}{q} - \eta(x - x_0)^2 \quad (10)$$

This function has local maxima at positions $x_0 \pm \frac{1}{2\eta q}$ (provided these points fall within Region 2) and a local minimum at x_0 .

Figure 2 illustrates this pattern of absolute advantage across the knowledge domain. Counter to their raw decision-making advantage, genius workers have their greatest advantage not at the furthest points from known knowledge (where their costs are highest) but at the boundaries where routine workers cannot function at all, but genius costs remain moderate.

4.2 Optimal Allocation Rule

Given this pattern of absolute advantage, we can now formulate a simple and intuitive allocation rule that solves the manager's problem:

Proposition 1 (Optimal Allocation Rule). *When routine workers are in sufficient supply (L_R is large enough to cover all questions in Region 2), the optimal allocation rule is:*

1. Rank all questions $x \in [0, 1]$ by their absolute advantage $AA(x) = V_G(x) - V_R(x)$
2. Allocate genius workers to the questions with the highest absolute advantage until the genius capacity L_G is exhausted
3. Allocate routine workers to all remaining questions where $V_R(x) > 0$
4. Leave unanswered any questions where no genius capacity remains and $V_R(x) = 0$

This allocation maximises the total value of answers provided.

Proof. The proof relies on a simple exchange argument. Suppose there exists an allocation where a question with lower absolute advantage $AA(x_1)$ is assigned to a genius worker, while a question with higher absolute advantage $AA(x_2) > AA(x_1)$ is not. Exchanging these assignments would increase total value by $AA(x_2) - AA(x_1) > 0$. Therefore, any optimal allocation must assign genius workers to questions in descending order of absolute advantage. $\square$

This allocation rule has an appealing economic interpretation: the manager should assign each type of worker to questions where they create the most value relative to the other type, respecting capacity constraints. When genius workers are scarce, they should be allocated to questions where their relative advantage is highest.

4.3 Equilibrium Allocation Patterns

The shape of the absolute-advantage function implies two regimes; see Figure 2. In Region 2, absolute advantage is $AA(d) = \frac{d}{q} - \eta d^2$, which peaks at $d^* = \min(q, \frac{1}{2\eta q})$. Thus, when $\eta \leq 1/(2q^2)$, scarce genius capacity is directed first to the edges of the routine domain, $x_0 \pm q$, where the incremental value over routine is greatest. When $\eta > 1/(2q^2)$, the peaks are interior: initial genius allocations go to narrow bands centred at $x_0 \pm d^*$, where the advantage over routine is locally maximised. In both cases, the remainder of Region 2 is served by routine workers, while Regions 1 and 3 are unanswered until genius capacity expands.

As L_G rises, these initial bands widen in the manner illustrated in Figure 2. In the edge-first regime, the expansion proceeds outward into Regions 1 and 3 and inward toward x_0 , with the centre left last because geniuses' relative advantage is smallest at $d = 0$. In the interior-first regime, the bands around $x_0 \pm d^*$ broaden symmetrically; once their marginal

advantage falls to the edge level $1 - \eta q^2$, additional capacity is placed at $x_0 \pm q$ and then extended further outward. At intermediate L_G , the allocation takes the form of two mirror-image strips flanking x_0 . With further increases, these strips grow toward one another, shrinking the central routine region; in the limit, abundant genius capacity covers the entire domain.

5 Transformational Impact of AI Geniuses

The emergence of transformative artificial intelligence (AI) systems capable of generating creative solutions and novel insights fundamentally alters the landscape of knowledge work allocation. As described by Amodei (2024), these systems can potentially provide “a country of geniuses in a data center,” dramatically expanding the supply of genius-level cognitive capacity. In this section, we extend our model to incorporate AI geniuses alongside human geniuses and routine workers, analysing how optimal allocation patterns and income distribution change in both the short and long run.

5.1 AI Geniuses: Capabilities and Cost Structure

We model AI geniuses as having fundamentally similar capabilities to human geniuses—they can create new knowledge for any question in the domain, including those where uncertainty exceeds the threshold for routine workers. However, we allow for potential differences in efficiency by introducing a distinct cost parameter.

Definition 1 (AI Genius Workers). *AI genius workers can generate exact answers for any question, with a cost structure similar to human geniuses but potentially less efficient:*

$$\text{Cost}_{AI}(x) = \eta_{AI} \cdot (x - x_0)^2 \quad (11)$$

where $\eta_{AI} \geq \eta$ represents the efficiency parameter for AI geniuses relative to human geniuses.

The net value created by an AI genius answering a question at position x is therefore:

$$V_{AI}(x) = 1 - \text{Cost}_{AI}(x) = 1 - \eta_{AI} \cdot (x - x_0)^2 \quad (12)$$

This value is depicted in Figure 3. The assumption that $\eta_{AI} \geq \eta$ reflects the possibility that AI systems, while capable of genius-level work, may require more resources or exhibit lower efficiency than top human geniuses when creating novel knowledge far from established

frameworks. This formulation creates a natural hierarchy: human geniuses have a comparative advantage over AI geniuses, who in turn have a comparative advantage over routine workers in regions where uncertainty is high.⁵

To analyse the impact of introducing AI geniuses into the knowledge economy, we distinguish between two time horizons:

- **Short run:** The existing allocation thresholds for routine workers remain fixed, reflecting organisational rigidities, contractual obligations, or adjustment costs
- **Long run:** The manager fully optimises worker allocation across all three worker types, potentially leading to significant structural changes in the knowledge economy

Crucially, we assume that human geniuses remain scarce ($L_G < 1$) but AI geniuses are available in effectively unlimited supply. This assumption reflects the scenario described by Amodei (2024), where millions of instances of AI models can be deployed once developed.

5.2 Short-Run Analysis: Fixed Routine Worker Allocation

In the short run, we assume that routine workers continue to operate within the thresholds established prior to the introduction of AI geniuses. Specifically, if routine workers were previously allocated to a set of questions A_R^* , this allocation remains unchanged immediately following the introduction of AI. This assumption captures institutional rigidities, existing contractual arrangements, and adjustment costs that prevent immediate reoptimisation of the entire knowledge workforce.

The manager’s immediate problem becomes reallocating human geniuses and allocating the newly available AI geniuses to maximise total value, taking the routine worker allocation as given:

$$\max_{A_G, A_{AI}} \int_{A_G} V_G(x) dx + \int_{A_{AI}} V_{AI}(x) dx \quad (13)$$

subject to capacity constraints $|A_G| \leq L_G$, non-overlap conditions $A_G \cap A_{AI} = \emptyset$, and respecting the fixed routine worker allocation A_R^* .

Proposition 2 (Short-Run Allocation with AI Geniuses). *Given fixed routine worker allocation A_R^* from the pre-AI equilibrium, the short-run optimal allocation with AI geniuses follows these principles:*

⁵If $\eta_{AI} < \eta$ then AI-genies would immediately substitute for all human geniuses regardless of the scarcity of the latter.

1. Human geniuses are reallocated to questions where their absolute advantage over AI geniuses is highest: $AA_{H/AI}(x) = V_G(x) - V_{AI}(x) = (\eta_{AI} - \eta) \cdot (x - x_0)^2$
2. Human geniuses are allocated to questions furthest from the known knowledge point within their previous allocation domain
3. AI geniuses replace human geniuses in positions closer to the known point where human geniuses were previously allocated
4. AI geniuses are allocated to previously unanswered questions where $V_{AI}(x) > 0$

This creates a distinctive three-tier allocation pattern where human geniuses handle the most distant questions, AI geniuses handle intermediate questions, and routine workers maintain their original allocation.

Proof. The absolute advantage of human geniuses over AI geniuses at position x is:

$$AA_{H/AI}(x) = V_G(x) - V_{AI}(x) \tag{14}$$

$$= (1 - \eta \cdot (x - x_0)^2) - (1 - \eta_{AI} \cdot (x - x_0)^2) \tag{15}$$

$$= (\eta_{AI} - \eta) \cdot (x - x_0)^2 \tag{16}$$

Since $\eta_{AI} \geq \eta$ by assumption, this advantage is non-negative and strictly increasing in $(x - x_0)^2$. Therefore, human geniuses have the greatest advantage over AI geniuses at positions furthest from the known knowledge point.

For a given pre-AI allocation A_G^* of human geniuses, the optimal reallocation involves: (i) Ranking positions in A_G^* by distance from the known point, $|x - x_0|$; (ii) Allocating human geniuses to positions with the highest $|x - x_0|$ values until the capacity L_G is exhausted; (iii) Allocating AI geniuses to the remaining positions in A_G^* where routine workers are not already allocated and (iv) Allocating AI geniuses to previously unanswered questions where $V_{AI}(x) > 0$. This allocation maximises total value because it preserves the total value from routine workers (which remains fixed in the short run) and it maximises the value from the combined human and AI genius allocation by assigning each type to positions where they have a comparative advantage. The three-tier pattern emerges naturally from the comparative advantages of each worker type. $\square$

The short-run allocation pattern has an interesting economic interpretation: the introduction of AI geniuses leads to a specialisation of human geniuses in the most novel and challenging questions, while AI geniuses take over questions that are moderately difficult; see Figure 3. Routine workers initially maintain their previous role, handling questions close to the known point.




Figure 3: Short-Run and Long-Run Worker Allocation with AI Geniuses

5.3 Long-Run Analysis: Optimised Worker Allocation

In the long run, the manager can fully reoptimise worker allocation across all three types, potentially leading to significant structural changes in the knowledge economy. The manager's long-run problem becomes:

$$\max_{A_R, A_G, A_{AI}} \int_{A_R} V_R(x) dx + \int_{A_G} V_G(x) dx + \int_{A_{AI}} V_{AI}(x) dx \quad (17)$$

subject to capacity constraints $|A_R| \leq L_R$ and $|A_G| \leq L_G$, and non-overlap conditions ($A_R \cap A_G = A_R \cap A_{AI} = A_G \cap A_{AI} = \emptyset$).

Proposition 3 (Long-Run Allocation with AI Geniuses). *In the long-run equilibrium with abundant AI genius capacity:*

1. Routine workers are completely displaced from knowledge work if η_{AI} is sufficiently close to η such that $V_{AI}(x) > V_R(x)$ for all x where $V_R(x) > 0$

2. Human geniuses are allocated exclusively to questions furthest from the known knowledge point, specifically:

$$A_G^{LR} = \{x \in [0, 1] : |x - x_0| \geq d^*\} \quad (18)$$

where d^* is the distance threshold such that $|A_G^{LR}| = L_G$

3. AI geniuses handle all remaining questions where either $V_{AI}(x) \geq V_R(x)$ or $V_R(x) = 0$

If η_{AI} is significantly larger than η such that $V_{AI}(x) < V_R(x)$ for some values of x , then routine workers continue to handle those questions in the long run.

Proof. For any question position x , the manager compares the values $V_R(x)$, $V_G(x)$, and $V_{AI}(x)$ and assigns the question to the worker type providing the highest value.

The comparative advantage of human geniuses over AI geniuses is $AA_{H/AI}(x) = (\eta_{AI} - \eta) \cdot (x - x_0)^2$, which is strictly increasing in $(x - x_0)^2$. Therefore, human geniuses have the greatest advantage over AI geniuses at positions furthest from the known knowledge point.

For positions where $V_R(x) > 0$, the advantage of AI geniuses over routine workers is:

$$AA_{AI/R}(x) = V_{AI}(x) - V_R(x) \quad (19)$$

$$= (1 - \eta_{AI} \cdot (x - x_0)^2) - \left(1 - \frac{|x - x_0|}{q}\right) \quad (20)$$

$$= \frac{|x - x_0|}{q} - \eta_{AI} \cdot (x - x_0)^2 \quad (21)$$

If η_{AI} is sufficiently close to η such that $AA_{AI/R}(x) > 0$ for all x where $V_R(x) > 0$, then AI geniuses completely displace routine workers. This condition is more likely to be satisfied when the efficiency gap between human and AI geniuses is small.

The allocation of human geniuses to questions furthest from the known point follows directly from maximising their comparative advantage over AI geniuses. The distance threshold d^* is determined by the capacity constraint $|A_G^{LR}| = L_G$.

If η_{AI} is significantly larger than η such that $AA_{AI/R}(x) < 0$ for some x , then routine workers retain a comparative advantage over AI geniuses for those questions and continue to handle them in the long run. $\square$

The long-run allocation leads to a fundamental reorganisation of knowledge work, as illustrated in Figure 3, with several key features. First, human geniuses become ultra-specialised in handling the most novel, challenging, and distant questions—those furthest from established knowledge. Second, previously unanswered questions (in Regions 1 and 3) now receive

answers from either human or AI geniuses, expanding the total domain of questions that can be addressed. Finally, once routing is re-optimised and AI genius capacity is available at scale, routine work is fully displaced unless AI is itself inefficient relative to the routine threshold. The displacement condition is $V_{AI}(d) \geq V_R(d)$ for all $d \in [0, q]$, i.e. $\frac{d}{q} - \eta_{AI}d^2 \geq 0$ on that interval, which holds iff $\eta_{AI} \leq 1/q^2$. When $\eta_{AI} \leq 1/q^2$, AI replaces routine throughout Region 2, Regions 1 and 3 are handled by geniuses, and scarce human geniuses specialise at the far frontier where their advantage over AI rises with distance.⁶

6 Conclusion

This paper has examined how transformative AI systems providing “genius on demand” would affect knowledge worker allocation and labour market outcomes. Our analysis reveals that when genius capacity is scarce, optimal allocation directs genius workers to questions where their absolute advantage over routine workers is highest — notably at domain boundaries where uncertainty is too high for routine workers. The introduction of AI geniuses transforms this allocation pattern, with human geniuses specialising in questions furthest from existing knowledge, where their comparative advantage over AI is greatest.

These findings contribute to our understanding of knowledge work allocation, but have modelled AI geniuses as independent agents that replace human workers. This approach, while analytically tractable, may not capture the most promising potential of AI technologies. There may be categories of geniuses, such as artisans, artists or persuasive geniuses. It is not clear if the Amodeli conception might cover those categories as well or if AI-genuses might also be more efficient at doing what we have modelled as routine work in this paper (see, for example, Gans (2025b) for an example of this type of AI). If there were AI agents that undertook knowledge inference, then it could easily be the case that in the short run, routine workers would be impacted, with geniuses perhaps being moved to more novel tasks in the long run. The point here is that being specific regarding how human and AI workers engage with knowledge will impact predictions regarding transformative AI.

Future research should examine how AI could function as tools in conjunction with knowledge workers rather than as automated independent agents (Agrawal et al., 2025). Such collaborative human-AI systems might leverage complementary strengths: human contextual understanding, ethical judgment, and creative insights alongside AI computational power and

⁶If $\eta_{AI} > 1/q^2$, routine workers survive only at the outer edge of their feasible set. The crossover solves $V_{AI}(d) = V_R(d)$, giving $d = 1/(\eta_{AI}q)$, so routine is used on $d \in (1/(\eta_{AI}q), q]$, AI handles the interior $d < 1/(\eta_{AI}q)$, and all $d > q$ go to geniuses. The measure actually assigned to routine in this edge band depends on the scarcity of human geniuses L_G : with tighter L_G , AI covers more of the interior, leaving routine only this outer ring; as $\eta_{AI} \downarrow 1/q^2$ the ring shrinks to zero width.

pattern recognition capabilities. Models of such complementary human-AI workflows would enrich our understanding of knowledge work evolution beyond displacement narratives.

The economic and social value generated by transformative AI may ultimately derive not from standalone genius capabilities but from novel workflows and institutions that effectively combine human and AI capabilities. By reframing the relationship between AI and knowledge workers from one of substitution to one of complementarity, we may discover approaches to knowledge creation that exceed what either humans or machines could achieve independently. The genius of the future may reside not exclusively in machines or humans, but in the symbiotic relationships between them.

References

- Acemoglu, Daron and Pascual Restrepo**, “The Race Between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment,” *American Economic Review*, 2018, *108* (6), 1488–1542.
- Agrawal, Ajay, Joshua S. Gans, and Avi Goldfarb**, “The Economics of Bicycles for the Mind,” *NBER Working Paper 30434*, 2025.
- Amodei, Dario**, “Machines of Loving Grace,” 2024. Available at <https://darioamodei.com/machines-of-loving-grace>.
- Carnehl, Christoph and Johannes Schneider**, “A Quest for Knowledge,” *Econometrica*, 2025, *93* (2), 623–659.
- Gans, Joshua S.**, “A Model of the Babbage Firm,” *NBER Working Paper*, 2025.
- , “A Quest for AI Knowledge,” *NBER Working Paper*, 2025.
- Garicano, Luis**, “Hierarchies and the Organization of Knowledge in Production,” *Journal of Political Economy*, 2000, *108* (5), 874–904.
- Halmos, Paul Richard**, “The legend of John von Neumann,” *The American Mathematical Monthly*, 1973, *80* (4), 382–394.
- Ide, Enrique and Eduard Talamas**, “Artificial Intelligence in the Knowledge Economy,” *Journal of Political Economy*, 2025. forthcoming.