The Information Value of Artificial Intelligence: Evidence from Firm Growth and Labor Allocation

Boyan Jovanovic* New York University

Sai Ma† Federal Reserve Board

This Version: November 2026

The latest version is available here.

Abstract

We develop an assignment model in which AI raises firm productivity through two channels: a novel information channel that provides better information for business decisions and a more conventional production channel that makes task execution more efficient. Empirically, we use earnings-call transcripts dated before ChatGPT's November 2022 release to measure each firm's predetermined exposure to both channels. From 2023 to 2025, firms one standard deviation higher in information channel exposure show a 1.62-percentage-point larger decline in management revenue-guidance error, about 20% of its pre-2023 mean. Disagreement among analysts about future earnings also declines significantly more for these firms. The same firms shift hiring toward information roles, such as data scientists and research analysts, offer these hires higher wages, and record 2.1 percentage points faster annual growth in sales per worker. Greater production channel exposure also predicts higher sales per worker, largely through lower employment and without improving guidance accuracy. The estimated model closely reproduces the shift toward information hiring and the faster growth in sales per worker associated with information exposure, without targeting either outcome. Under the baseline calibration, the model attributes about half of the post-2022 rise in sales per worker to the information channel.

JEL Classification: O33, J24, J31, D22, D83

Keywords: artificial intelligence, signal precision, growth, assignment, worker-firm sorting

*Department of Economics, NYU, 19 W.4th St, 6th Floor, New York, NY 10012. (boyan.jovanovic@nyu.edu)

†Federal Reserve Board of Governors, C Ave & 20th Street NW, Washington, DC 20551. (sai.ma@frb.gov). We thank seminar participants at Princeton University and the Federal Reserve Board for helpful comments and Jack Cook and Martina Cuneo for excellent research assistance. The views expressed are those of the authors and not necessarily those of the Federal Reserve Board or the Federal Reserve System.

1 Introduction

Information is a key input into production. Firms hire workers and select projects before fully observing demand, project quality, customer risk, or the fit of a worker to a task. The noisier the signals about these conditions, the more mistakes firms make and the lower the return to information skill. Workers with those skills may also fail to reach the firms where they are most productive. AI can raise productivity through two distinct channels: better information about "what to do" and greater efficiency in "how to do it." The *information channel* makes the firm's signals more precise, whereas the *production channel* reduces the labor and time needed to carry out the chosen task. More precise signals raise output through better decisions, worker-firm assignment, and the accumulation of capabilities inside the firm.

Existing measures of firms' exposure to large language models (LLMs) do not distinguish more efficient production from better information. For example, Eloundou, Manning, Mishkin, and Rock (2024) and Eisfeldt, Schubert, Zhang, and Taska (2026) score how much generative AI tools can speed up the tasks in each occupation. These scores describe the work a firm's employees perform but not how much more precise information would improve the firm's decisions. Our contribution is to measure each firm's exposure to both channels and to connect these exposures, through a model of signal precision and assignment, to the accuracy of management's revenue guidance, hiring, and output.

To distinguish the two channels, we develop a multidimensional assignment model following Lindenlaub (2017). Firms differ in their information and execution capabilities, while workers differ in the corresponding skills. Each firm hires one worker, uses a noisy signal to decide what to do, and then carries out that decision. Decision efficiency is the expected accuracy of that choice. More precise information raises decision efficiency and the return to information skill, and it changes worker-firm assignment. Greater production efficiency instead raises output from carrying out the chosen task without improving the decision itself. We embed this assignment problem in an overlapping-generations framework in which the worker assigned to a firm also contributes to its future capabilities, a contribution we call learning from matched workers.

The model yields three predictions that distinguish the channels. First, once AI becomes available, firms more exposed to the information channel make better decisions, whereas firms more exposed to the production channel do not. Second, better information raises the return to information skill, shifts worker-firm assignment toward information-skilled workers,

and raises their wages. Third, both channels can raise output. Holding assignment fixed, the output gain from the information channel is larger when the firm's initial signal is less precise.

ChatGPT's public release in November 2022 provides a common date at which generative AI tools became widely accessible. We study public firms using two measures fixed before this date and built from the firms' own earnings calls, in which managers describe the decisions they face and how the firm produces. First, **information channel exposure**, or information exposure, ranks firms by how much language-model extraction, comparison, and synthesis of existing text and records could improve the precision of the estimates behind their decisions, holding production technology fixed. Second, **production channel exposure**, or production exposure, ranks firms by how much generative AI tools can reduce the labor or time needed to carry out a chosen task, holding information precision fixed. Three accompanying scores measure how uncertain the estimates behind firm decisions are, how fragmented and hard to compare their inputs are, and how much is at stake in those decisions. Because each firm receives both exposure scores, we compare post-2022 changes across firms that differ in one score while holding the other fixed.

The return to better information should also depend on how much the firm knew initially, because a given increase in precision is more valuable when initial precision is low. We measure this lack of initial precision, which we call opacity, with disagreement, range, and realized error in analysts' pre-2022 earnings forecasts from I/B/E/S. In the accounting and finance literature, these forecast properties are standard measures of the quality of the firm's information environment and residual uncertainty about future earnings.[1]

The post-2022 changes in guidance, hiring, and firm performance support all three predictions. First, firms with greater information exposure issue more accurate revenue guidance. A one-standard-deviation difference in information exposure is associated with a 1.62-percentage-point larger decline in the absolute error of management's revenue guidance, about one-fifth of its mean before 2023. Management guidance also becomes more accurate relative to a benchmark forecast built from the firm's and its industry's sales histories. Because the benchmark's own error does not change, the improvement is consistent with better

[1]Lang and Lundholm (1996) relate disclosure quality to analyst following and the properties of analysts' earnings forecasts. Barron, Kim, Lim, and Stevens (1998) show how forecast dispersion, consensus error, and the number of forecasts can be used to infer properties of analysts' information environment. Zhang (2006) uses analyst coverage and forecast dispersion as proxies for information uncertainty, and Diether, Malloy, and Scherbina (2002) use analyst forecast dispersion to measure differences of opinion.

information available to management. Disagreement among analysts about future earnings also declines for firms with greater information exposure. By contrast, production exposure does not predict guidance accuracy.

Second, labor demand shifts toward the tasks that better information makes more valuable. From Capital IQ professional histories, we find that the share of new hires in information roles, such as data scientists and research analysts, rises significantly by 10.5 percentage points per standard deviation of information exposure, relative to a pre-2023 mean of 20.6 percent. The shift reflects both more information hires and fewer hires into execution roles, such as software developers and engineers, and it begins in 2023. The same firms file more H-1B applications for information roles and offer higher wages for these positions within occupation.

Third, firms with greater information exposure grow faster in sales per worker, a revenue-based measure of labor productivity. From 2023 to 2025, annual growth is 2.07 percentage points faster per standard deviation of information exposure, after removing each firm's 2018–2022 trend. On the same sample, sales rise and employment does not change, so the faster growth does not come from cutting headcount. Information exposure is associated with even faster growth among firms with greater initial opacity, in line with the model's third prediction. Production exposure is associated with 1.53 percentage points faster annual sales-per-worker growth, but through lower employment and without any change in guidance accuracy. These different patterns for sales and employment match the model's distinction between an information channel operating through decisions and a production channel that saves labor.

The timing and cross-sectional tests distinguish information exposure from pandemic recovery, industry conditions, and general AI exposure. Guidance accuracy, information hiring, and sales-per-worker growth begin to change only in 2023, with no differential trend in earlier years, and placebo dates produce no comparable change in guidance error or sales-per-worker growth. The guidance and sales-per-worker estimates also remain statistically significant when we absorb more detailed industry effects and control for pandemic-era volatility, overhiring, and prior guidance errors. In the cross section, the uncertainty, fragmentation, and stakes scores do not reproduce the guidance or sales-per-worker findings, and controlling for the workforce exposure measure of Eisfeldt, Schubert, Zhang, and Taska (2026) leaves the estimates for information exposure essentially unchanged.

Information exposure also predicts later generative AI (GenAI) use in independently scored financial reports outside the earnings calls used to construct the measure. The improvement in guidance tracks this use: the decline in guidance error associated with information exposure is larger among firms that use GenAI in more types of financial reports and has no counterpart under a 2021 placebo date.

We next ask whether the assignment model can jointly account for the evidence on decisions, workers, wages, and firm performance. We estimate the model using three groups of moments: errors in management guidance, pre-2023 worker sorting and the wages of information hires, and the 2023–2025 sales-per-worker and employment responses to production exposure. Neither the post-2022 hiring response nor the firm-performance responses to information exposure enter estimation. Without these responses, the model predicts a 10.8-percentage-point increase in the share of new hires in information roles, compared with 10.5 in the data, and 2.11 percentage points faster annual sales-per-worker growth, compared with 2.07 in the data.

Under the baseline calibration, the estimated model attributes about half of the post-2022 rise in sales per worker to the information channel. That calibration translates gains in decision efficiency into output using the share of tasks that Eloundou, Manning, Mishkin, and Rock (2024) estimate language models can improve. Because their measure encompasses both channels, we also consider smaller task shares for decision efficiency. The information channel's contribution remains substantial under these alternatives, although the part due to reassigning workers is more sensitive to how assignment is computed.

After a one-time rise in precision, learning from matched workers raises output until firms settle at a higher level, without a permanent increase in the growth rate. To obtain sustained growth, we allow firms to invest in more signals. Greater signal precision raises the return to this investment, increasing long-run growth and expanding the upper tail of the distribution of firm growth.

Taken together, the joint pattern across guidance, labor allocation, and firm performance distinguishes the information and production channels.

Related literature. The assignment model builds on work in which workers and firms have several characteristics. Lindenlaub (2017) develops a quadratic-Gaussian structure in which these characteristics jointly determine sorting. We use this structure to distinguish firm capabilities and worker skills in information and execution. Lise and Postel-Vinay (2020)

add search frictions and skill accumulation to multidimensional sorting, while Edmond and Mongey (2026) study skill prices when workers supply several skills as a bundle. Lindenlaub, Oh, Rodriguez, and Veldkamp (2026) show that AI adoption depends on comparative advantage relative to workers' skills rather than exposure alone. Sorensen (2007) uses a two-sided matching model to distinguish sorting from the influence of investors on the firms they finance. That distinction parallels our separation of the output gain at the existing assignment from the gain due to reassigning workers.

Learning gives assignment a dynamic role. In Jovanovic (2014), more accurate signals about worker skill improve matching and shape the human capital of the next generation. Our model instead considers signals about the problems a firm and worker must solve together. Jovanovic and Rousseau (2026) interpret AI as an increase in the precision of such signals. Building on the information-investment mechanism in Jovanovic and Ma (2023), we show that more precise signals raise the return to investing in more signals.

The labor-market implications relate to employer learning and technological change. Farber and Gibbons (1996), Altonji and Pierret (2001), and Lange (2007) show that wages and assignments come to reflect worker skill as employers learn over time. Autor, Chin, Salomons, and Seegmiller (2026) show that newly created, technology-linked work initially commands a wage premium, and Hassan, Kalyani, and Restrepo (2026) study how scarce expertise in new technologies raises the skill premium until that expertise diffuses. In our model, better information raises the return to information skill, with a larger wage response for workers who have more of that skill. Because workers have separate information and execution skills, AI can change both their relative returns and their assignment across firms.

Our empirical analysis also relates to work on AI exposure and firm performance. Babina, Fedyk, He, and Hodson (2024) find that firms investing in AI grow through product expansion. Eisfeldt, Schubert, Zhang, and Taska (2026) score occupational tasks with a language model, weight the scores by each firm's occupational employment, and show that workforce exposure to generative AI is capitalized in equity values. Benmelech and Zator (2022) examine automation and robots, Czarnitzki, Fernandez, and Rammer (2022) relate AI to measured productivity, and Gao (2026) studies AI and the allocation of effort. At the worker level, Brynjolfsson, Li, and Raymond (2023) document large productivity gains for workers given access to AI tools. By contrast, Humlum and Vestergaard (2025) find little average effect on earnings or recorded hours across eleven AI-exposed Danish occupations during

the first two years after ChatGPT's release, and Acemoglu (2025) argues that the aggregate productivity gains implied by task-level cost savings are modest. These worker-level and task-based studies focus on productivity or labor savings in exposed tasks. The information channel instead raises the precision of the information behind firm decisions, and firms more exposed to it raise sales without changing employment. In the spirit of the text-based firm measures of Hassan, Hollander, van Lent, and Tahoun (2019), our exposure measures come from firm text rather than occupational tasks, so the two channels can differ within the same firm.

Constructing firm measures from earnings calls connects the paper to recent work that uses large language models to extract economic information from text. Jha, Qian, Weber, and Yang (2024) and Kim, Muhn, and Nikolaev (2023) measure corporate policies and risks from earnings calls. Iacoviello and Tong (2026) and Hartley (2025) measure geopolitical and economic-policy risk, Bybee (2024) recovers beliefs from news, and Bhandari, Jugulum, Veldkamp, and Venkateswaran (2026) measure firm scope from annual reports. Ash, Hansen, and Muvdi (2024) survey this emerging literature. Recent econometric work by Ludwig, Mullainathan, and Rambachan (2026), Carlson and Dell (2025), and Hansen (2026) studies validation and inference when language-model measures are used in empirical work. Validation and inference both matter in our setting because a language model constructs firm characteristics from text dated before ChatGPT became available, and we relate those characteristics to post-2022 changes in guidance, hiring, and productivity.

Our use of management and analyst forecasts also connects the information channel to work on forecasting under imperfect information. Grossman and Stiglitz (1980) show how signal precision is capitalized in prices. Whereas informed investors observe the same signal in their model, Albagli, Hellwig, and Tsyvinski (2023) allow informed investors to receive heterogeneous signals, which generates disagreement. Closer to our measurement, Bali, Kelly, Moerke, and Rahman (2023) use machine-learning forecasts to measure disagreement, and Jansen, Nagel, Yannelis, and Zhang (2022) study how better data improve screening in credit markets. Bonelli (2026) finds that venture capitalists adopting data technologies screen firms resembling historical comparables more accurately while financing fewer novel startups. Available data are therefore more valuable when they are more comparable, a property captured by our information fragmentation score. Building on this work, we measure decision efficiency using the accuracy of management's own revenue guidance and initial opacity using

analyst disagreement and forecast error.

The remainder of the paper proceeds as follows. Section 2 develops the model and its predictions. Section 3 describes the data and constructs the exposure measures. Section 4 presents the empirical tests. Section 5 estimates the model and decomposes the cumulative 2023–2025 increase in sales per worker, then studies long-run growth when firms invest in more signals. Section 6 reports robustness and discusses interpretation. Section 7 concludes. The Online Appendix provides data and specification details, proofs, and supporting results.

2 A Model of Signal Precision and Assignment

We study a multidimensional assignment model following Lindenlaub (2017), while allowing firms to make decisions under uncertainty. Each match performs two tasks. In the information task, the firm uses a noisy signal to choose its decision, and in the execution task, it carries out that decision. Firms differ in their capabilities in the two tasks, while workers differ in the corresponding skills, and firms and workers match competitively. AI can raise the signal's precision or the efficiency with which the decision is carried out. Because the young worker contributes to the firm's capabilities, today's assignment also affects future output.

2.1 Environment

Agents – Time is discrete, with two-period overlapping generations. At date t, a unit measure of old firms forms one-period matches with a unit measure of young workers. Firm i has an execution (E) capability and an information (I) capability, collected in the vector $X^F_{i,t} = (X^F_{E,i,t}, X^F_{I,i,t})'$. Across firms at date t, these capability vectors have distribution ν^F_t, with mean μ^F_t and covariance $\Sigma_{F,t}$. Firm i hires one young worker, whose execution and information skills are similarly collected in $X^W_{i,t} = (X^W_{E,i,t}, X^W_{I,i,t})'$. Across workers, these skill vectors have distribution ν^W, with mean μ^W and covariance Σ_W. The distribution of worker skills is fixed over time, whereas firm capabilities evolve with assignment. The old firm operates for one period and consumes its profits. The young worker takes over firm i when she becomes old, so their match shapes the capabilities with which she operates the firm in the next period.

Information – Firm capabilities and worker skills are known when the match forms, but the payoff-relevant state θ is not. The firm's decision d, such as the scale of an investment or

the choice of an operating technology, must reflect this state. As in Jovanovic and Rousseau (2026), AI improves the precision of a signal about this state. For firm i,

$$\theta \sim \mathcal{N}\left(\bar{\theta}, \frac{1}{p_0}\right), \qquad s_i = \theta + \xi_i, \qquad \xi_i \sim \mathcal{N}\left(0, \frac{1}{p_{A,i}}\right),$$

where $\bar{\theta}$ and p_0 are the common prior mean and precision, the noise ξ_i is independent of θ, and $p_{A,i}$ is the precision of the firm's signal. AI reduces signal noise and thereby raises $p_{A,i}$. We measure decision efficiency relative to prior uncertainty,

$$e(d) = 1 - p_0(\theta - d)^2,$$

so a decision made on the prior alone has zero expected efficiency. The firm maximizes expected efficiency by choosing $d_i^* = E[\theta \mid s_i]$. Normal updating gives expected decision efficiency

$$q_i \equiv q(p_{A,i}) = E[e(d_i^*)] = \frac{p_{A,i}}{p_0 + p_{A,i}}, \qquad q'(p_{A,i}) > 0, \qquad q''(p_{A,i}) < 0. \tag{1}$$

Expected decision efficiency is the share of prior uncertainty that the signal resolves, and it rises with precision at a decreasing rate. Realized efficiency varies around q_i with variance

$$\text{var}[e(d_i^*)] = 2(1 - q_i)^2,$$

so greater precision raises expected decision efficiency and reduces its dispersion.

Match output – Signal precision determines how accurately the firm chooses what to do. Generative AI can also make it more efficient to carry out that choice. We denote this increase in production efficiency by g_i and set it to zero before generative AI is available.[2] To combine these two channels with worker-firm assignment, we follow Lindenlaub (2017) and treat firm capabilities and worker skills as complementary inputs into match output. Because each firm hires one worker, match output is also output per worker. Suppressing time subscripts, its expected value is

$$Q_i = \bar{Q} + a_{q,i} q_i + g_i + (X_i^F)' A(q_i, g_i) X_i^W, \qquad A(q_i, g_i) = \begin{pmatrix} (1+g_i)a_E & c \\ c & q_i a_I \end{pmatrix}, \tag{2}$$

where $\bar{Q}$ is output common to all matches and $a_{q,i} \geq 0$ measures how much output depends directly on firm i's decision efficiency. The matrix $A(q_i, g_i)$ determines how firm capabilities

[2] Online Appendix B.2 provides a microfoundation and derives $g_i = \lambda_i/(1-\lambda_i)$, where λ_i is the share of execution operations that generative AI tools perform without labor.

and worker skills combine. On its diagonal, $a_E > 0$ and $a_I > 0$ govern the complementarity between firm capability and worker skill within the execution and information tasks.[3] Off the diagonal, c links firm capability in one task with worker skill in the other, making them complements when $c > 0$ and substitutes when $c < 0$.

Information and production channels – Equation (2) shows how AI raises output through two channels. Through the *information channel*, a more precise signal raises expected decision efficiency q_i and makes firm information capability and worker information skill more productive together. Through the *production channel*, a rise in g_i raises output directly and increases the marginal product of execution skill at firms with greater execution capability.

Decision efficiency and output – Holding the worker fixed, the information channel raises output through the two terms of equation (2) that contain decision efficiency q_i, the direct decision component $a_{q,i}q_i$ and the information interaction $q_i a_I X^F_{I,i} X^W_{I,i}$. The marginal contribution of decision efficiency to output is therefore

$$\ell_i \equiv \frac{\partial Q_i}{\partial q_i} = a_{q,i} + a_I X^F_{I,i} X^W_{I,i}. \tag{3}$$

Thus more precise information raises output when $\ell_i > 0$, and it raises output more at firms that have greater information capability and hire workers with greater information skill.

Equation (2) describes expected output because assignment occurs before the state is realized. Once the state is realized, $e(d_i^*)$ replaces q_i in the direct decision component and the information interaction. Taking expectations over the decision error recovers equation (2) because $E[e(d_i^*)] = q_i$.

Law of motion for firm capabilities – The current match also affects future production because the young worker contributes to firm i's capabilities in the next period. Firm capability levels evolve according to

$$X^F_{i,t+1} = b + \Phi X^F_{i,t} + \Gamma X^W_{i,t} + \varepsilon_{i,t+1}, \tag{4}$$

where $b \in \mathbb{R}^2$ is a constant vector, the 2×2 matrix Φ governs how firm capabilities carry over from one period to the next, the 2×2 matrix Γ measures how much the young worker's skills contribute to next period's firm capabilities, and the shock $\varepsilon_{i,t+1}$ is independent and Gaussian with mean zero and covariance Σ_ε. The shock is observed at the beginning of period $t+1$. We assume that $\|\Phi\| < 1$, so firm capabilities would be stable without the

[3] Online Appendix B.3 shows how Cobb–Douglas production within each task yields the two diagonal interactions in equation (2). The empirical implementation does not impose the Cobb–Douglas restrictions.

contribution of worker skills through Γ. We call this contribution learning from matched workers. Averaging equation (4) across matches gives

$$\mu_{t+1}^F = b + \Phi\mu_t^F + \Gamma\mu^W. \tag{5}$$

Subtracting this mean equation from equation (4) and dropping the firm subscript gives the law of motion for firm capabilities relative to their means,

$$x_{t+1}^F = \Phi x_t^F + \Gamma x_t^W + \varepsilon_{t+1}, \tag{6}$$

where $x_t^F = X_t^F - \mu_t^F$ and $x_t^W = X_t^W - \mu^W$ are the deviations from the means.

Timing – At the beginning of period t, old firms have capabilities X^F given by equation (4), and young workers draw skills X^W from ν^W. The market observes firm capabilities and worker skills and forms matches for current production. Each firm then observes its signal, chooses d, and produces output. Wages are paid, and equations (5) and (6) determine next period's mean and distribution of firm capabilities.

2.2 Competitive Equilibrium

An assignment π_t pairs firm types (X^F, p_A, g, a_q) with worker skill vectors X^W. The wage of worker type X^W is $W_t(X^W)$. Because an old firm operates for one period, it chooses its worker to maximize current profits,

$$\max_{X^W} \left\{ Q_i(X^F, X^W, p_A, g, a_q) - W_t(X^W) \right\}. \tag{7}$$

The worker's contribution to firm i's capabilities in the next period, given by equation (4), therefore does not enter this choice. A worker accepts the highest wage offered for her type.[4]

Definition 1 *A competitive equilibrium is a sequence* $\{\nu_t^F, \pi_t, W_t\}_{t\geq 0}$ *such that: (i) almost every pair in* π_t *solves equation* (7)*, (ii) each worker accepts the highest available wage, (iii)* π_t *has the given firm and worker marginals, and (iv)* ν_{t+1}^F *is generated by equation* (4) *under* π_t*.*

[4] In Jovanovic (2014), the young worker also values the expected profit of the firm she operates in the next period. In Online Appendix C.2, we extend the model so that the worker's payoff from a match also includes this profit and show that the estimates remain highly robust.

2.3 Analytical Results

Equilibrium assignment – The direct terms $a_{q,i}q_i$ and g_i in equation (2) do not vary with the worker whom firm i hires. Assignment instead depends on how much the firm values each worker skill. Because only the term $(X_i^F)'A(q_i, g_i)X_i^W$ depends on the worker, the firm's marginal values of execution and information skill are

$$\frac{\partial Q_i}{\partial X_{E,i}^W} = (1+g_i)a_E X_{E,i}^F + cX_{I,i}^F, \qquad \frac{\partial Q_i}{\partial X_{I,i}^W} = cX_{E,i}^F + q_i a_I X_{I,i}^F. \tag{8}$$

Thus an increase in production efficiency g_i raises the marginal value of execution skill most at firms with greater execution capability. Greater signal precision, by contrast, works through decision efficiency q_i and raises the marginal value of information skill most at firms with greater information capability. In vector form, the marginal values are

$$\underbrace{m_{i,t}^F}_{2\times 1} \equiv \left(\frac{\partial Q_i}{\partial X_{E,i}^W}, \frac{\partial Q_i}{\partial X_{I,i}^W}\right)' = \underbrace{A(q_i, g_i)'}_{2\times 2}\underbrace{X_{i,t}^F}_{2\times 1}, \tag{9}$$

so two firms with the same capabilities but different q_i or g_i value worker skills differently and therefore hire different workers. Across firms, $m_{i,t}^F$ has mean $\mu_{m,t}$ and covariance $\Sigma_{m,t}$.

As in Lindenlaub (2017), two conditions in Definition 1 determine the assignment. First, market clearing, condition (iii), requires that every worker be hired, so the skills of hired workers have the distribution ν^W. Second, firm optimality, condition (i), rules out pairs of firms that could produce more by exchanging workers. Because such an exchange between firms i and j changes their combined output by $-(m_i^F - m_j^F)'(X_i^W - X_j^W)$, the assignment must be positively sorted in marginal values,

$$\left(m_i^F - m_j^F\right)'\left(X_i^W - X_j^W\right) > 0. \tag{10}$$

Theorem 1 (Competitive equilibrium) *Suppose $m_{i,t}^F$ and $X_{i,t}^W$ are Gaussian with non-singular covariance matrices $\Sigma_{m,t}$ and Σ_W. Then the unique equilibrium assignment is*

$$X_{i,t}^W = \mu^W + B_t\left(m_{i,t}^F - \mu_{m,t}\right), \tag{11}$$

where

$$B_t = \Sigma_{m,t}^{-1/2}\left[\Sigma_{m,t}^{1/2}\Sigma_W\Sigma_{m,t}^{1/2}\right]^{1/2}\Sigma_{m,t}^{-1/2}. \tag{12}$$

The assignment preserves the worker distribution and is positively sorted in firms' marginal values of worker skills. A quadratic wage schedule supports this allocation.

Proof. See Online Appendix B.

With one skill, as in the log-normal case of Jovanovic (2014), B_t would be the standard deviation of worker skill divided by the standard deviation of marginal values. A firm with a marginal value one standard deviation above the mean would then hire a worker with skill one standard deviation above the mean. With two skills, equation (12) generalizes this ratio to the symmetric positive definite matrix that satisfies $B_t\Sigma_{m,t}B_t = \Sigma_W$, so the hired workers have covariance Σ_W.

When decision efficiency q and production efficiency g are common across firms, $m_t^F - \mu_{m,t} = A(q,g)'x_t^F$, where x_t^F and x_t^W are the deviations from the means in equation (6). The assignment in equation (11) then gives the hired worker's skills directly from the firm's capabilities,

$$x_t^W = T_t(q,g)x_t^F, \qquad T_t(q,g) = B_tA(q,g)'.$$

If $A(q,g)$, $\Sigma_{F,t}$, and Σ_W are diagonal with positive diagonal entries, $T_t(q,g)$ is diagonal with positive entries, so firms and workers are positively sorted within each task. Otherwise, sorting is positive in m_i^F but need not be positive within each task.

Because hired workers contribute to next period's firm capabilities, assignment also shapes their distribution. Under common q and g, substituting $x_t^W = T_t(q,g)x_t^F$ into equation (6) gives

$$x_{t+1}^F = M_t(q,g)x_t^F + \varepsilon_{t+1}, \qquad M_t(q,g) = \Phi + \Gamma T_t(q,g),$$

so $M_t(q,g)$ combines how firm capabilities carry over, through Φ, with learning from matched workers, through $\Gamma T_t(q,g)$. In a stationary distribution of firm capabilities, these matrices are constant, and we write them as $T(q,g)$ and $M(q,g)$. Provided every eigenvalue of $M(q,g)$ has modulus less than one, the stationary covariance of firm capabilities satisfies

$$\Sigma_F = M(q,g)\Sigma_F M(q,g)' + \Sigma_\varepsilon. \tag{13}$$

Setting $\mu_{t+1}^F = \mu_t^F$ in equation (5) gives the stationary mean

$$\mu^F = (I_2 - \Phi)^{-1}(b + \Gamma\mu^W), \tag{14}$$

where I_2 is the 2×2 identity matrix. The stationary mean depends on average worker skill but not on assignment. Assignment instead affects the covariance of firm capabilities. Because the worker distribution is fixed, it does so only through the covariance between firm capabilities and hired workers' skills, $\Sigma_{F,t}T_t(q,g)'$.

Information, assignment, and output – A firm-specific increase in signal precision $p_{A,i}$ raises output holding assignment fixed and also changes the worker the firm hires. Proposition 1 separates these two effects. Because one firm has measure zero, its increase leaves the wage schedule and the matrix B_t unchanged.

Proposition 1 (Information raises output) *Consider an increase in firm i's signal precision at the prevailing wage schedule. Holding its current capabilities fixed, the assigned worker changes according to*

$$\frac{dX_i^W}{dp_{A,i}} = B_t \begin{pmatrix} 0 \\ a_I X_{I,i}^F q'(p_{A,i}) \end{pmatrix}. \tag{15}$$

Current output satisfies

$$\frac{dQ_i}{dp_{A,i}} = \underbrace{\ell_i q'(p_{A,i})}_{\text{assignment held fixed}} + \underbrace{(m_i^F)' B_t \begin{pmatrix} 0 \\ a_I X_{I,i}^F q'(p_{A,i}) \end{pmatrix}}_{\text{worker reassignment}}. \tag{16}$$

The reassignment also changes next period's firm capabilities by Γ times the right-hand side of equation (15). If $A(q_i, g_i)$, B_t, and Γ are diagonal with positive entries, $X_{I,i}^F > 0$, and $\ell_i > 0$, both output terms are positive and next period's information capability rises.

Proof. See Online Appendix B.

The first term in equation (16) holds the worker fixed. The second reflects the output produced by the worker whom the firm hires after its demand for information skill rises. That worker then raises firm i's information capability in the next period through equation (4). The diagonal case gives the signs in closed form, while the estimated model allows the cross-task interaction c to affect both skills of the assigned worker.

Holding assignment fixed, the output gain from greater signal precision also depends on how precise the firm's initial signal is. Because decision efficiency $q(p_A)$ is concave, the same increase in precision raises q_i by more when initial precision $p_{A,i}$ is lower, given the common prior precision p_0. We call less precise initial signals initial opacity.

Corollary 1 (Initial opacity and output) *Hold assignment fixed and suppose $\ell_i > 0$. Then*

$$\left.\frac{\partial Q_i}{\partial p_{A,i}}\right|_{X_i^W} = \ell_i \frac{p_0}{(p_0 + p_{A,i})^2} > 0, \qquad \left.\frac{\partial^2 Q_i}{\partial p_{A,i}^2}\right|_{X_i^W} < 0. \tag{17}$$

For the same increase in precision, the output gain is larger at a firm with lower initial signal precision.

Proof. See Online Appendix B.

Proposition 1 concerns a firm-specific increase in precision, which changes assignment. To isolate the wage response, we next consider a market-wide increase. We restrict assignment to the information dimension and let firms begin with common signal precision, so a common increase strengthens the complementarity between firm information capability and worker information skill at every firm while leaving the assignment unchanged. Proposition 2 gives the wage schedule under this restriction of equation (2).

Proposition 2 (Information-skill wages) *Hold execution capabilities and skills fixed and set $c = 0$. Suppose prior and signal precision are common across firms and the information-capability and information-skill distributions imply $X_I^W = \kappa X_I^F$ in equilibrium, where $\kappa > 0$. The wage of a worker with information skill X_I^W is*

$$W(X_I^W, q) = \bar{W} + \frac{a_I q}{2\kappa}(X_I^W)^2. \tag{18}$$

The intercept $\bar{W}$ is the outside option of a worker with zero information skill. For $X_I^W > 0$, $dW(X_I^W, q)/dp_A > 0$.

Proof. See Online Appendix B.

Because the assignment does not change, firms bid more for the same information-skilled workers, and the higher complementarity shows up in wages.

Corollary 2 (AI raises the information-skill premium)
Under Proposition 2, if $X_I^W > 0$,

$$\frac{\partial^2 W}{\partial p_A \, \partial X_I^W} = \frac{a_I}{\kappa} q'(p_A) X_I^W > 0.$$

Thus AI raises wages more for workers with greater information skill.

Proof. See Online Appendix B.

The production channel instead raises production efficiency g_i in execution tasks without changing decision efficiency q_i. Holding assignment fixed, Proposition 3 shows that higher production efficiency raises both output and the marginal value of worker execution skill.

Proposition 3 (Production efficiency) *Hold signal precision and assignment fixed. If firm execution capability and worker execution skill are positive, then*

$$\frac{\partial Q_i}{\partial g_i} = 1 + a_E X_{E,i}^F X_{E,i}^W > 0, \qquad \frac{\partial^2 Q_i}{\partial g_i \, \partial X_{E,i}^W} = a_E X_{E,i}^F > 0. \tag{19}$$

Thus an increase in production efficiency raises output and the firm's marginal value of worker execution skill.

Proof. See Online Appendix B.

Forecast errors – Output alone does not distinguish the two channels, because both decision efficiency q and production efficiency g raise output Q. Errors in the firm's forecast of its own output separate them. The firm forecasts the expected output in equation (2), but because it chooses d before θ is observed, realized decision efficiency e can differ from the expected efficiency $q(p_A)$ on which the match was formed, and realized output can then differ from the forecast. Production efficiency g, by contrast, does not affect this forecast error. We collect forecast errors unrelated to the decision in $u_i \sim \mathcal{N}(0, \sigma_u^2)$, independent of the state and the signal. Firm i's resulting forecast error in output is $\ell_i[e(d_i^*) - q_i] + u_i$, which has mean zero, and Lemma 1 gives its dispersion as a function of signal precision.

Lemma 1 (Signal precision and forecast errors) *Conditional on firm i's match, the variance of the forecast error is*

$$\operatorname{var}\left(\ell_i[e(d_i^*) - q_i] + u_i \mid X_i^F, X_i^W\right) = 2\ell_i^2(1 - q_i)^2 + \sigma_u^2. \tag{20}$$

Forecast-error variance decreases with signal precision whenever $\ell_i \neq 0$. For a fixed increase in precision, the decline is larger when initial signal precision is lower or $|\ell_i|$ is larger. When $\sigma_u^2 = 0$, mean absolute forecast error is

$$4\varphi(1)|\ell_i|(1 - q_i),$$

where φ is the standard-normal density.

Proof. See Online Appendix B.

A more precise signal therefore reduces the forecast error created by the decision, especially when the firm begins with less precise information or when decision efficiency contributes more to output.

Over time, a persistent firm-specific increase in $p_{A,i}$ changes not only the worker the firm hires but also, through learning from matched workers, its information capability. A firm-specific increase in production efficiency g_i likewise changes assignment along the execution dimension. After a one-time increase in precision, however, output converges to a higher level, so growth rises only while firm capabilities and worker-firm assignment adjust. Section 5.5 therefore allows firms to invest in more signals, which can sustain long-run growth.

3 Data and Measurement

The information and production channels have different implications for decisions, worker-firm assignment and wages, and firm performance. To study these differences, we measure firms' exposure to each channel from earnings calls dated before ChatGPT's release and combine the transcript measures with management and analyst forecasts, worker records, and accounting data.[5]

3.1 LLM-Based Firm Measures

We use Capital IQ earnings conference call transcripts from 2019 through November 10, 2022, accessed through Datastream, to characterize the information and production problems described by each firm. For each document, GPT-5.5 receives an excerpt with the firm name and ticker removed. To make the scores comparable across documents, we hold the classification prompt fixed, provide no examples of scored text, and set the temperature to zero. To keep the measures predetermined, GPT-5.5 receives no text dated after ChatGPT's release, no analyst forecasts, and no realized firm outcomes. Online Appendix Section A.1 reproduces the classification prompt and reporting format.

The classifier assigns each document five scores from 0 to 100, with higher values indicating more of the corresponding attribute. The first three characterize the information problem the firm potentially faces. The last two measure its exposure to generative AI through the information and production channels:

- *State uncertainty* (U_i) measures how much the firm must estimate payoff-relevant states, such as demand, project quality, or customer risk, before committing resources.
- *Information fragmentation* (F_i) measures whether the inputs used in those estimates are dispersed, qualitative, unstructured, or difficult to compare.
- *Decision stakes* (K_i) measures the economic importance and reversibility of the decisions that rely on those estimates.

[5] We harmonize all sources to Compustat's GVKEY. Capital IQ transcripts and professional histories are linked through Capital IQ company identifiers to GVKEY. Transcript links are checked against date-valid nine-digit CUSIPs, CRSP PERMNOs, and CRSP/Compustat Merged links. LSEG guidance is matched to I/B/E/S actuals primarily by eight-digit CUSIP, with official and current tickers used as fallbacks, and I/B/E/S securities are mapped from PERMNO to GVKEY using date-valid CRSP/Compustat Merged links. H-1B records are linked through a validated crosswalk from harmonized employer names to Compustat firms. We manually review conflicting or one-to-many links.

- *Information channel exposure* (I_i) measures how much language-model extraction, comparison, synthesis, or interpretation can improve the precision of those estimates, holding product demand and production technology fixed.
- *Production channel exposure* (P_i) measures how much generative AI tools can reduce the labor or time required to produce the firm's goods or services, holding information precision fixed.

We require at least four firm-quarters and retain documents for which the classifier's confidence, its 0 to 100 rating of whether the excerpt suffices for classification, is at least 70. We also exclude documents that describe only how hard the firm is to understand from the outside, without a corresponding internal information problem. We first average each score across eligible documents within a firm-quarter and then average the quarter means, giving each quarter equal weight. Under these restrictions, 4,477 firms contribute 104,594 earnings-call documents, 23 per firm on average. Table 1 reports the distribution of the five scores.

Table 1: Firm Measures: Summary Statistics

Variable	Firms	Mean	SD	p10	p50	p90
State uncertainty	4,477	49.97	7.86	40.86	49.20	59.65
Information fragmentation	4,477	49.94	7.64	39.97	49.99	59.53
Decision stakes	4,477	49.94	7.66	40.13	49.84	59.74
Information channel exposure, I	4,477	49.64	17.75	30.98	49.05	70.53
Production channel exposure, P	4,477	50.00	13.99	32.31	49.42	68.62

Notes. Scores range from 0 to 100 and use 104,594 earnings-call documents dated through November 10, 2022.

Two channels within firms. The measures describe two channels through which AI can affect the same firm, not mutually exclusive types of firms, so every firm receives both scores. Information channel exposure is high when language models make the estimates behind a decision more precise, whereas production channel exposure is high when generative AI tools reduce the labor or time required to carry out the chosen task. The two scores have a correlation of 0.18, and the four groups defined by whether each score is above or below its median each contain about one-quarter of the sample. Table 2 illustrates the distinction. For example, Shutterstock, which licenses stock images, footage, and music, has high exposure to both channels. Interpreting media and customer intent guides sourcing, search, and

Table 2: Examples of Information and Production Channel Exposure

Firm	Information Channel	Production Channel	I	P
Shutterstock	Interpreting diverse media and customer intent guides content sourcing, search relevance, and licensing.	Images, footage, music, designs, and custom creative content are the firm's products.	100	75
Burlington Stores	Buyers combine fashion trends, sales forecasts, vendor availability, and local demand to choose and allocate inventory.	Stores, distribution centers, and merchandise handling account for most production activity.	71	27
Gartner	Credit assessments, client retention, acquisitions, and tax positions require judgment from firm records.	Research, consulting, and conferences are the firm's main outputs.	28	70
IDEXX Laboratories	The firm forecasts diagnostic demand, inventory needs, supply disruptions, and uncertain R&D outcomes.	It produces diagnostic products, laboratory services, software, and imaging systems.	1	35

Notes. The table selects one illustrative firm from each of the four combinations formed by splitting I and P at their sample medians. Selection does not use outcomes measured after 2022. Scores come from earnings calls held before ChatGPT's release and range from 0 to 100. Descriptions are based on the firms' independently classified 2021–2022 Form 10-K filings.

licensing decisions, while producing media and custom content is itself part of the firm's output. The off-price retailer Burlington Stores is more exposed to the information channel through its inventory decisions, where language models can refine demand and availability estimates, than through its physical retail and distribution operations. By contrast, the research and advisory firm Gartner is more exposed to the production channel, because its research and consulting are outputs sold to clients rather than information used primarily in its own decisions. Finally, the veterinary diagnostics company IDEXX Laboratories is below the median on both measures. Its earnings calls describe uncertain diagnostic demand and R&D outcomes but give little indication that extracting, comparing, or synthesizing existing text and records would make the estimates behind its decisions materially more precise.

The distinction also appears within industries. Differences across four-digit industries account for only 12.8 percent of firm-level variation in information channel exposure, compared with 46.9 percent of the variation in production channel exposure, so most of the variation in information channel exposure lies within industries. Online Appendix Table A2 pairs firms in the same detailed industry with similar production channel exposure. In each pair, independently scored 10-K filings give the same ordering of information channel exposure.

3.2 Other Data: Forecasts, Labor, and Firm Performance

Management forecasts come from LSEG Company Guidance. For annual revenue guidance that can be compared with realized Compustat revenue, we record three properties of the forecast. Absolute midpoint error is the absolute difference between the midpoint of management's range and realized revenue, expressed as a percentage of realized revenue. Miss distance records how far the realization falls outside the range and equals zero when it is inside, while range width records how tightly management states its forecast. We assign each observation to the calendar quarter in which management issues the guidance rather than to the fiscal year it covers. Annual guidance is issued 250 to 430 days before the fiscal year end, most often in the first quarter. In the model, the resulting forecast errors reflect uncertainty that remains unresolved when management forms its revenue forecast and fall as signal precision rises.

To distinguish better management information from firms becoming mechanically easier to forecast, we also construct two forecasts available before the realization. The firm-history benchmark extrapolates the firm's trailing three-year sales growth, whereas the shrinkage benchmark places equal weight on the firm and its two-digit industry's trailing growth. For either benchmark, management's advantage is benchmark error minus management error. The benchmark's own error then provides a placebo: if I_i merely selects firms that became easier to forecast, both management and the benchmark should improve.

We measure hiring and wages in information and execution roles using two sources. Labor Condition Application records from the Department of Labor cover H-1B, H-1B1, and E-3 workers. Among AI- and machine-learning-related applications, we distinguish information from execution roles using job titles and occupational codes fixed before examining firm outcomes. Information roles include data scientists, statisticians, operations-research and business-intelligence analysts, economists, and related decision-support occupations, whereas execution roles include software developers, systems engineers, architects, database administrators, and engineering occupations. Certified wages then measure the price of information skills within occupation, year, and state, and the same filings measure AI hiring before ChatGPT's release, which we call prior AI hiring. Capital IQ professional histories provide the second, broader measure of labor allocation. We classify job titles as information or execution roles using the dictionary described in Online Appendix Section A.2 and measure the share of new hires in information roles. A new hire is a person's first observed role at the

firm, and we retain one role per person-firm spell.

Compustat provides annual sales, employment, cost of goods sold, and operating expenses. Following Babina, Fedyk, He, and Hodson (2024), we measure labor productivity with sales per worker, a revenue-based measure.[6] Sales growth and employment growth are constructed on the identical sample, so their coefficients satisfy the accounting identity that sales-per-worker growth equals sales growth minus employment growth.

I/B/E/S provides firm-level analyst disagreement, forecast range, absolute forecast error, and analyst coverage, the number of analysts following the firm. Analyst disagreement, forecast range, and forecast error form the opacity measure O_i, built from data before 2022. The opacity measure is constructed independently of the transcript-based state uncertainty score U_i and determines initial precision when we estimate the model. In Section 4.6, we later use its quarterly version after 2022 to ask whether firms with greater information channel exposure became easier to understand from the outside. To validate the transcript measures against characteristics measured before ChatGPT's release, we also use the industry AI exposure measure of Felten, Raj, and Seamans (2021) and the workforce exposure measure of Eisfeldt, Schubert, Zhang, and Taska (2026).

3.3 Validation and Measurement Concerns

Before turning to the empirical results, we address the two concerns that Ludwig, Mullainathan, and Rambachan (2026) raise for economic applications of language models: validation of constructed measures and inadvertent use of future information in prediction. The language model is used here for measurement, but a recently trained language model may have learned the later outcomes of the firms that the measures rank. Section 3.3.1 examines whether the scores reflect later events using alternative models, a different source document, and human assessments. Section 3.3.2 compares them with firm characteristics measured independently before 2023.

3.3.1 LLM Look-Ahead Bias

A recent language model may know how generative AI subsequently spread through firms. Were the model to lean on this knowledge, the scores could reflect how firms later fared rather than the problems they described at the time, a concern Iacoviello and Tong (2026)

[6] Like other revenue-based measures, sales per worker can also reflect prices, markups, product mix, and outsourcing. The regressions absorb common industry movements with two-digit industry-by-year effects.

raise for historical text scored by the same class of models. Recent evidence documents this risk in financial and macroeconomic applications.[7] The prompt limits direct reliance on such knowledge. It removes the firm name and ticker, supplies only the excerpt, and asks the model to score the operating decisions the text describes rather than to identify the company or predict its later success. Removing the identifiers also makes it harder to draw on general knowledge about a named firm, although contextual clues may still identify the firm, as documented by Lopez-Lira, Tang, and Zhu (2025) and Engelberg, Manela, Mullins, and Vulicevic (2025). Wu, Yang, Ying, and Zhou (2025) show that more aggressive anonymization can itself remove economically meaningful content from earnings calls. Telling a model to set aside what it knows does not erase that knowledge from its parameters, so we examine the scores in three further ways.

First, we ask whether the rankings survive changes in model developer, size, and vintage. Claude Opus 4.8 and Claude Sonnet 5 come from a different developer, while Claude 3 Haiku and the open-weight Llama 3 8B were released in the first half of 2024 and therefore leave a shorter window for information from after ChatGPT's release to enter their training data. All four models receive the same prompt. Panel A of Online Appendix Table A1 shows that the information channel exposure scores from the four alternatives correlate between 0.82 and 0.91 with the baseline scores. The corresponding correlations for production channel exposure are 0.93 with Haiku and 0.91 with Llama. The rankings are therefore stable across models of different vintages and sizes, including proprietary and open-weight families.[8] Because Llama's weights are publicly available, its classifications also provide an open-weight version of the measures that can be reconstructed outside a proprietary system.

Second, we score a different type of document. Earnings calls and annual reports are written for different audiences and serve different purposes. Calls are conversations with analysts about operations and outlook, while 10-K filings are regulatory documents organized around audited financial statements, business descriptions, and risk disclosures. Following the use of 10-K text to measure firm characteristics in Bhandari, Jugulum, Veldkamp, and Venkateswaran (2026), we apply the same definitions using GPT-5.5 to the 2021–2022 10-K filings of firms in the management-guidance sample, without firm names, tickers, the

[7] See Glasserman and Lin (2023), Sarkar and Vafa (2025), Lopez-Lira, Tang, and Zhu (2025), Crane, Karra, and Soto (2025), and Levy (2026).

[8] He, Lv, Manela, and Wu (2025) train models only on text available by a fixed date, which provides a stricter test. Engelberg, Manela, Mullins, and Vulicevic (2025) propose iterative entity removal, and Gao, Jiang, and Yan (2025) test statistically whether passages appeared in training data.

earnings-call scores, or any outcome data. For the 389 firms scored from both documents, the rank correlation between the two sources is 0.86 for production channel exposure, with a 95 percent bootstrap interval of 0.83 to 0.89. Combining state uncertainty, information fragmentation, decision stakes, and information channel exposure gives a correlation of 0.50, with an interval of 0.41 to 0.58. For information channel exposure alone, the correlation is 0.40, with an interval of 0.30 to 0.49. After accounting for production channel exposure, firm size, and detailed industry in each source, the rank correlation for information channel exposure remains 0.33, with an interval of 0.17 to 0.46. Production activities are described in similar terms in both documents, which helps explain the stronger cross-document agreement for production channel exposure. A firm's information problems instead surface mainly in the discussion of operations and outlook that calls emphasize. The cross-document agreement indicates that the measures capture persistent firm characteristics rather than the wording of particular calls. The conditional correlation also shows that the agreement in information channel exposure is not inherited from production activities or industry. Panel B of Online Appendix Table A1 reports these correlations, and Online Appendix Figure A1 shows the corresponding score differences across earnings-call quintiles.

Third, we compare the language model's classifications with human judgment. Two research assistants independently scored information channel exposure and production channel exposure for 100 randomly selected firms using the same rubric and scoring range, without seeing the language-model scores. Panel C of Online Appendix Table A1 shows that their correlations with the language-model scores range from 0.87 to 0.89 for information channel exposure and from 0.79 to 0.91 for production channel exposure. The close agreement supports treating the language-model scores as reliable proxies for human classifications of firm disclosures.

3.3.2 Predetermined External Validation

If the transcript scores capture firms' information and production problems, they should be related to predetermined firm characteristics measured independently of the transcripts. Comparing the transcript scores with these predetermined characteristics parallels the external validation of language-model measures in Jha, Qian, Weber, and Yang (2024) and Kim, Muhn, and Nikolaev (2023).

Table 3 compares the transcript measures with independently measured proxies for firm opacity, operating volatility, and AI exposure. Panel A regresses each characteristic on state

uncertainty U, I, and P jointly, except that the regression for AI hiring includes only I and P. Across these specifications, I is positively related to analyst disagreement and forecast range at the 1 percent level and to operating volatility and prior AI hiring at the 10 percent level. In the disagreement regression, however, the coefficient on U is about twelve times that on I. The larger coefficient on U is consistent with U capturing the level of uncertainty and I capturing how much language models can improve the precision of the firm's estimates, which we call language-model fit.

Language-model fit is only one determinant of the output gain from greater precision. Equation (1) implies that the same increase in precision is more valuable when initial precision is lower. State uncertainty, which moves closely with analyst disagreement in Panel A, provides a transcript measure of low initial precision. The gain also depends on how fragmented the inputs to the firm's estimates are and on the stakes of the decisions that rely on them. We therefore supplement I_i with adjusted information channel exposure, defined as

$$I_i^{adj} = (U_i F_i K_i I_i)^{1/4}. \tag{21}$$

The geometric mean gives the four scores equal weight and is small when any one of them is small. The adjusted measure I_i^{adj} correlates 0.78 with I_i. Online Appendix Table A1 shows that, as for I_i, versions constructed from Haiku and Llama scores correlate 0.86 and 0.88 with the baseline adjusted measure.

Panel B compares I^{adj} with P across the same predetermined firm characteristics. Adjusted information channel exposure is positively related to analyst disagreement, forecast range, and absolute forecast error, and is unrelated to industry AI exposure. The production channel exposure measure is instead related to prior AI hiring and industry AI exposure and is unrelated to forecast range and absolute forecast error. A Wald test rejects equal coefficients for every characteristic in the panel. Together, these coefficients distinguish firms with greater adjusted information channel exposure from firms whose production was already exposed to AI before ChatGPT's release. Although I alone is only weakly related to prior AI hiring, I^{adj} is negatively related to it, so firms with greater adjusted information channel exposure were less likely to have hired AI workers before November 2022.

The workforce exposure measure of Eisfeldt, Schubert, Zhang, and Taska (2026) provides a further test of whether the two transcript measures capture distinct dimensions of exposure. They score occupational tasks with a language model and weight the resulting

Table 3: Validation of Firm Measures against Predetermined Characteristics

Panel A. Predetermined firm characteristics

	Analyst disagreement (1)	Forecast range (2)	Operating volatility (3)	AI hiring (4)
State uncertainty	0.661***	0.562***	0.113***	–
	(0.021)	(0.021)	(0.015)	–
Information channel exposure, I	0.054***	0.059***	0.026*	0.011*
	(0.012)	(0.014)	(0.014)	(0.007)
Observations	4,245	4,476	3,720	4,477
Log assets control	Yes	Yes	Yes	Yes
Two-digit industry FE	Yes	Yes	Yes	Yes

Panel B. The two channels

	Analyst disagreement (1)	Forecast range (2)	Absolute forecast error (3)	AI hiring (4)	Industry AI exposure (5)
Adjusted information channel exposure, I^{adj}	0.158***	0.139***	0.080***	−0.016**	0.001
	(0.019)	(0.019)	(0.022)	(0.006)	(0.014)
Production channel exposure, P	−0.040**	0.004	−0.002	0.031***	0.350***
	(0.018)	(0.018)	(0.020)	(0.005)	(0.018)
p-value: $I^{adj} = P$	< 0.001	< 0.001	0.016	< 0.001	< 0.001
Observations	4,245	4,476	4,449	4,477	2,384
Log assets control	Yes	Yes	Yes	Yes	Yes
Two-digit industry FE	Yes	Yes	Yes	Yes	Yes

Notes. Dependent variables and transcript scores are standardized. Panel A includes U, I, and P, with P omitted from the table and the U coefficient omitted in column (4). Panel B includes I^{adj} and P. All regressions control for log assets and two-digit industry fixed effects. Heteroskedasticity-robust standard errors are in parentheses. Industry AI exposure is from Felten, Raj, and Seamans (2021). $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

occupational exposures by each firm's employment shares in March 2022. If production channel exposure captures how much generative AI tools can reduce the labor needed in production, it should be related to workforce exposure, but information channel exposure, which concerns the firm's information problems, need not be. Panel A of Online Appendix Table A3 confirms this prediction. With two-digit industry effects, a one-standard-deviation increase in production channel exposure is associated with a 0.14-standard-deviation higher workforce exposure, and the coefficient on information channel exposure is small and statistically insignificant. Within four-digit industries, the coefficient on production channel exposure remains positive at 0.064, while the coefficient on information channel exposure is −0.007. The two coefficients differ with a p-value of 0.023. Production channel exposure is therefore related to the established measure of task exposure even within detailed indus-

tries, whereas information channel exposure captures a distinct characteristic of the firm's information problems.

4 Empirical Tests and Results

The baseline design uses the November 2022 release of ChatGPT as a common technology break. For outcome Y_{it}, we estimate

$$Y_{it} = \alpha_i + \delta_{s(i)t} + \beta_I Post_t \times I_i + \beta_P Post_t \times P_i + \text{controls}_{it} + \varepsilon_{it}, \tag{22}$$

where $Post_t$ equals one from 2023 onward and zero before, and β_I and β_P are the coefficients on its interactions with information and production channel exposure. The term α_i denotes firm fixed effects, $\delta_{s(i)t}$ denotes the time effects used for outcome Y_{it}, and $s(i)$ denotes the industry or guidance type relevant to that outcome. Any additional controls enter through controls_{it}, and ε_{it} is the error term. In the regressions, I_i and P_i are standardized once across the 4,477 scored firms. Both interactions enter jointly in every baseline specification. Because I_i and P_i use only text dated before ChatGPT's release, the coefficients compare changes after the common technology break across firms with different exposure to the information and production channels. The identifying assumption is that, conditional on the fixed effects and treatment of pre-2023 trends used for each dependent variable, these firms would not otherwise have experienced differential changes after 2022. The annual coefficients in Figure 1 provide the first test of this assumption. The model's predictions for guidance accuracy, labor allocation, and firm performance organize the evidence.

4.1 Management Guidance Accuracy

The first prediction is that better information improves decisions, and Lemma 1 shows that the same improvement lowers forecast errors. We test this prediction in Table 4 using direct measures of management revenue-guidance error and management's advantage over the two benchmark forecasts. From 2023 onward, a one-standard-deviation increase in information channel exposure is associated with a 1.617 percentage-point decline in absolute midpoint error, about one-fifth of its pre-2023 mean of 7.864 percentage points, and a 1.524 percentage-point decline in miss distance. The same firms also improve by 1.371 percentage points relative to their own historical benchmark and by 1.406 percentage points relative to the shrinkage benchmark.

Because the benchmark's own error does not improve with I_i, the improvement in management guidance goes beyond changes captured by the firm's and industry's sales history. By contrast, production channel exposure is unrelated to management-guidance error, and the equality tests reject a common I and P coefficient for absolute error, miss distance, and both benchmark advantages. Taken together, the direct guidance estimates, the two benchmark tests, and the absence of a relation with P_i support the information-channel interpretation.

The decline in guidance error does not reflect changes in forecastability or selection into guidance. As Online Appendix Table A9 reports, the guidance coefficient on information channel exposure is little changed when we absorb quarterly changes in forecastability within industries: interacting two-digit industry effects with the quarter of guidance issuance gives a coefficient of -1.621, and controlling continuously for forecast horizon or retaining only initial annual guidance gives -1.626 and -1.612. As for selection, information channel exposure predicts neither the decision to issue annual guidance nor the number of guidance observations.

Table 4: GenAI and Management Guidance Accuracy

	Absolute error (1)	Miss distance (2)	Range width (3)	History advantage (4)	Shrinkage advantage (5)	Shrinkage error (6)
Post $\times I$	-1.617^{***}	-1.524^{***}	-0.074	1.371^{***}	1.406^{***}	-0.234
	(0.271)	(0.255)	(0.053)	(0.413)	(0.403)	(0.332)
Post $\times P$	-0.051	-0.032	-0.082	-0.469	-0.732	-0.776
	(0.282)	(0.275)	(0.068)	(0.480)	(0.443)	(0.600)
p-value: $I = P$	< 0.001	< 0.001	0.935	0.005	0.001	0.476
Pre-2023 mean	7.864	6.568	2.766	8.635	5.979	13.856
Observations	5,041	5,041	5,041	4,910	4,910	4,910
Firms	869	869	869	859	859	859
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Guidance type $\times$ quarter FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes. The sample consists of annual revenue guidance. Errors and advantages are in percentage points. Miss distance is realized revenue outside the guidance range. Advantage is benchmark error minus management-guidance error. The firm-history benchmark uses trailing three-year sales growth, and the shrinkage benchmark averages firm and industry histories. Guidance type is defined by forecast horizon, point versus range guidance, and whether the observation updates earlier guidance. All columns include $Post \times I$, $Post \times P$, firm FE, and guidance type $\times$ quarter FE. Standard errors are two-way clustered by firm and quarter. Post begins in 2023. $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

4.2 Information Labor Demand, Hiring, and Wages

The second prediction concerns labor demand. Proposition 2 and Corollary 2 show that a market-wide increase in precision steepens the equilibrium return to information skill. When precision gains differ across firms, Proposition 1 implies that firms with larger precision gains demand and hire workers with more information skill. We therefore test whether firms with greater information channel exposure hire more workers into information roles and offer these hires higher wages within occupation. Table 5 examines hiring and wages using two independent data sources.

Panel A first uses Capital IQ professional histories to measure both the composition and the number of new hires. The share of new hires in information roles rises by 0.105 per standard deviation of I_i, a 51 percent increase relative to its pre-2023 mean of 0.206. The corresponding hire-count estimates show a parallel change in hiring levels.[9] Hires into information roles rise by about 50 percent, whereas hires into execution roles decline by about 14 percent. The composition and count estimates therefore both point to a shift toward information roles. The annual estimates in Figure 1 show that this shift begins in 2023.

Panel B turns to the narrower H-1B sample of AI- and machine-learning-related applications. Applications for information roles rise by 0.119 in $\log(1 + \text{applications})$ per standard deviation of I_i, while applications for execution roles change by 0.007 and are statistically indistinguishable from zero. At the same time, offered wages for information hires rise by 0.014 log units, or 1.4 percent, within occupation, state, and year. The increase in applications for information roles and the higher wages offered to these hires together support the rise in demand for information skill predicted by the model.[10] Production channel exposure instead predicts more Capital IQ hires into execution roles and fewer H-1B applications for information roles. It does not significantly predict the share of Capital IQ new hires in information roles, H-1B applications for execution roles, or information-hire wages.

[9] Columns (2) and (3) in Panel A of Table 5 estimate conditional Poisson regressions of annual hire counts with firm and two-digit industry-by-year effects to account for firm-years with zero hires. Their coefficients are proportional effects, so $\exp(\beta) - 1$ is the percentage change in expected hires. Conditioning on the firm effect drops firms with no hires of a given type during the sample. Online Appendix Section B.1 describes the specification. Online Appendix Table A13 reports $\log(1 + \text{hires})$ counts, which give smaller coefficients because most firm-years have zero or one classified hire.

[10] The wage estimate can reflect both the greater information skill of workers hired by high-exposure firms within an occupation and the higher equilibrium return to that skill. The model counterpart of this coefficient, which we use in estimation in Section 5.3, captures both, because it evaluates the change in the part of $W_t(X^W)$ that depends on assignment at the worker assigned to each firm.

Table 5: Information Labor Demand, Hiring, and Wages

Panel A. Capital IQ professional histories

	Share of new hires in information roles (1)	Information-role hires (2)	Execution-role hires (3)
Post $\times I$	0.105***	0.404***	−0.151***
	(0.013)	(0.070)	(0.041)
Post $\times P$	−0.021	−0.083	0.182**
	(0.017)	(0.139)	(0.071)
Pre-2023 mean	0.206	–	–
Observations	1,756	3,096	8,848
Firms	635	387	1,106

Panel B. H-1B AI/ML applications

	Information-role applications (1)	Execution-role applications (2)	Information-hire log wage (3)
Post $\times I$	0.119***	0.007	0.014***
	(0.020)	(0.022)	(0.005)
Post $\times P$	−0.114***	−0.072	−0.007
	(0.043)	(0.050)	(0.009)
Observations	1,824	1,824	25,970
Firms	228	228	223

Notes. In Panel A, a new hire is the first observed role at the firm. Column (1) reports the share of new hires in information roles. Columns (2) and (3) are conditional Poisson estimates of hire counts with firm and two-digit industry-by-year effects. They retain zero-hire years for firms with a positive count in the corresponding category. Panel B uses H-1B AI/ML applications. Columns (1) and (2) report log(1+applications). Column (3) reports log offered wages for information hires and includes state and occupation effects. All columns include $Post \times I$, $Post \times P$, and firm and year effects. Standard errors are clustered by firm. Post begins in 2023. $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

4.3 Firm Performance and Channel Separation

The third prediction, from Propositions 1 and 3, is that both channels can raise productivity. Panel A of Table 6 shows that firms with greater information channel exposure grow faster in sales per worker from 2023 onward, and the coefficient changes little with the treatment of earlier trends.

The sales-per-worker growth coefficient on information channel exposure in Panel A is 1.922 percentage points in the unadjusted specification, 2.074 when the projected pretrend is removed, and 1.910 with firm-specific trends. We use the pretrend-adjusted specification as the baseline because it removes differences in firms' growth before 2023. Moreover, the sales-per-worker, sales, and employment regressions use the same sample and transformation, so their coefficients satisfy the accounting identity exactly.

The production channel exposure measure also predicts faster sales-per-worker growth,

Table 6: Information Channel Exposure, Production Channel Exposure, and Firm Performance

Panel A. Two-channel estimates

	Sales-per-worker growth			Sales growth	Employment growth	Gross margin
	Unadjusted	Baseline	Firm trend	Pretrend adjusted	Pretrend adjusted	
	(1)	(2)	(3)	(4)	(5)	(6)
Post $\times I$	1.922***	2.074***	1.910***	2.053***	−0.021	1.264***
	(0.240)	(0.501)	(0.441)	(0.547)	(0.355)	(0.364)
Post $\times P$	1.420***	1.528**	1.274*	0.536	−0.992**	2.812***
	(0.330)	(0.721)	(0.669)	(0.746)	(0.477)	(0.451)
p-value: $I = P$	0.270	0.565	0.458	0.132	0.136	0.010
Observations	24,060	22,236	22,947	22,236	22,236	19,615
Firms	3,649	3,109	3,210	3,109	3,109	2,954
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Two-digit industry × year FE	Yes	Yes	Yes	Yes	Yes	Yes

Panel B. Opacity interactions and prior AI hiring

	Baseline (pretrend adjusted)	$I\times$ opacity	F-K-I composite × opacity	Add AI hiring
	(1)	(2)	(3)	(4)
Post $\times I$	2.074***	2.251***	–	1.973***
	(0.501)	(0.533)		(0.509)
Post × F-K-I composite	–	–	1.903***	–
			(0.555)	
Post $\times P$	1.528**	1.572**	1.707**	1.492**
	(0.721)	(0.716)	(0.720)	(0.724)
Post × opacity	–	−0.860	−0.701	–
		(0.762)	(0.752)	
Post $\times I\times$ opacity	–	1.834***	–	–
		(0.557)		
Post × F-K-I composite × opacity	–	–	1.314**	–
			(0.615)	
Post × AI hiring × opacity	–	–	–	0.257
				(2.763)
Observations	22,236	22,236	22,236	22,236
Firms	3,109	3,109	3,109	3,109
Firm and two-digit industry × year FE	Yes	Yes	Yes	Yes

Notes. The sample consists of Compustat firm-years from 2018 through 2025. Growth rates are winsorized and reported in percentage points. Gross margin is sales less cost of goods sold, divided by sales, and is also reported in percentage points. Pretrend adjustment removes each firm's projected 2018–2022 trend, while firm-trend specifications include firm-specific linear trends. Panel B uses pretrend-adjusted sales-per-worker growth. The F-K-I composite is the geometric mean of information fragmentation, decision stakes, and information channel exposure, expressed in standard-deviation units. Opacity averages pre-2022 analyst disagreement, range, and error after each is expressed in standard-deviation units. AI hiring indicates any AI- or machine-learning-related H-1B application filed before ChatGPT's release. All regressions include firm and two-digit industry-by-year fixed effects. Panel B, column (4), also includes the Post main effects for opacity and AI hiring. Standard errors are clustered by firm. $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

and the equality tests do not distinguish its coefficient from the coefficient on I_i. The two coefficients nevertheless differ in composition. For I_i, faster sales growth accounts for 2.053 of the 2.074-percentage-point coefficient, and lower employment accounts for the remaining 0.021. The employment estimate is near zero, with a 95 percent interval of about seven-tenths of a point on either side. The sales-per-worker growth coefficient on I_i therefore comes primarily from higher sales, consistent with the finding of Babina, Fedyk, He, and Hodson (2024) that firms investing in AI grow through product expansion. For P_i, the split reverses: lower employment accounts for about two-thirds of its 1.528-percentage-point coefficient, and the sales component of 0.536 is not distinguishable from zero. The point estimates imply different decompositions, although the equality tests for sales and employment have p-values of 0.132 and 0.136.

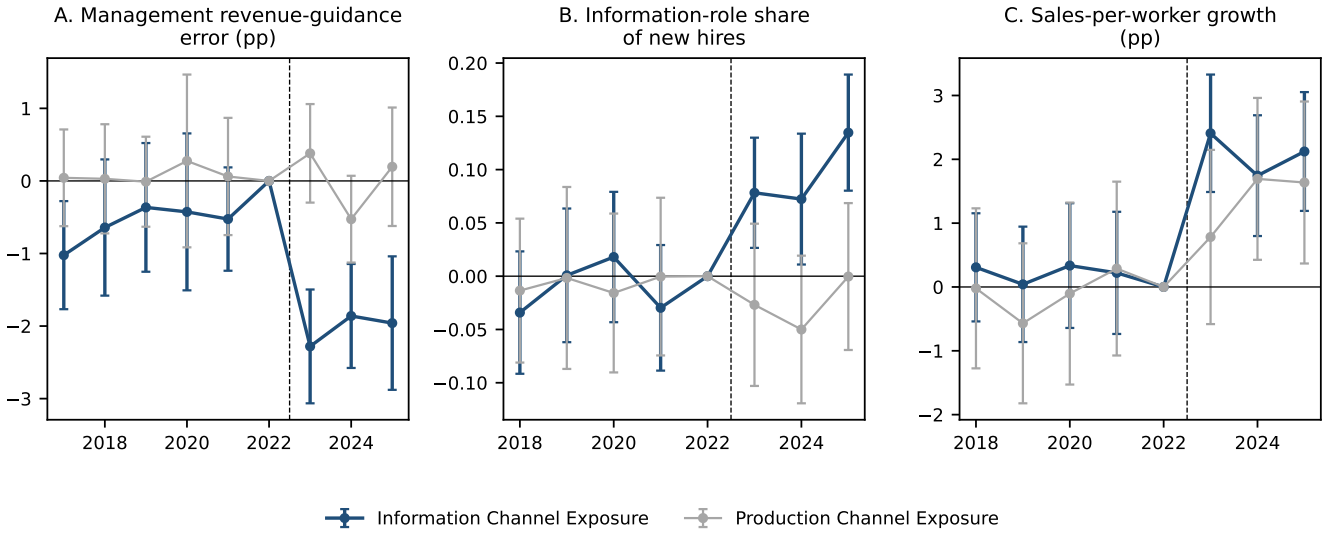
Guidance accuracy improves only with I_i, and margins further distinguish the channels. Gross margins rise by 1.26 and 2.81 percentage points per standard deviation of I and P, respectively, from a pre-2023 mean of 34.9 percent. The increase is more than twice as large for P_i, and the equality test rejects a common coefficient with a p-value of 0.010. Online Appendix Table A14 similarly shows that operating income per worker rises significantly only with P_i. The margin and profitability results therefore lend greater support to the production channel, whereas guidance accuracy distinguishes the information channel. Because margins can also vary with prices or product mix, they do not by themselves identify the channel.

Within the information channel, firms whose information tasks are better suited to language models grow faster in sales per worker when initial opacity is greater. Corollary 1 predicts this pattern for the output gain holding assignment fixed, and Panel B of Table 6 shows that it also holds for total sales-per-worker growth. In column (2), the interaction between I_i and analyst-based opacity is positive. Column (3) also yields a positive interaction with opacity when I_i is replaced by the F-K-I composite, the geometric mean of information fragmentation, decision stakes, and information channel exposure expressed in standard-deviation units. The composite leaves out state uncertainty because opacity enters separately as the measure of initial precision.

4.4 Timing of the Responses

The annual coefficients in Figure 1 show when guidance error, information hiring, and sales-per-worker growth begin to change. In Panel A, the guidance coefficients before 2022 are

Figure 1: Timing of Guidance Accuracy, Information Hiring, and Sales-per-Worker Growth



Notes. The figure reports year-specific coefficients on information channel exposure I and production channel exposure P, with 95 percent confidence intervals. The omitted year is 2022. The dashed line marks the November 2022 release of ChatGPT. Panel A uses management revenue-guidance error, Panel B the share of new hires in information roles from Capital IQ professional histories (Table 5, column (1), with firm and year FE), and Panel C unadjusted sales-per-worker growth, with firm and two-digit industry-by-year fixed effects. A new hire is a person's first observed role at the firm. Standard errors are clustered by firm. Joint p-values for the I coefficients before 2022 are 0.19, 0.42, and 0.93 in Panels A–C. The tested years are 2017–2021 in Panel A and 2018–2021 in Panels B–C.

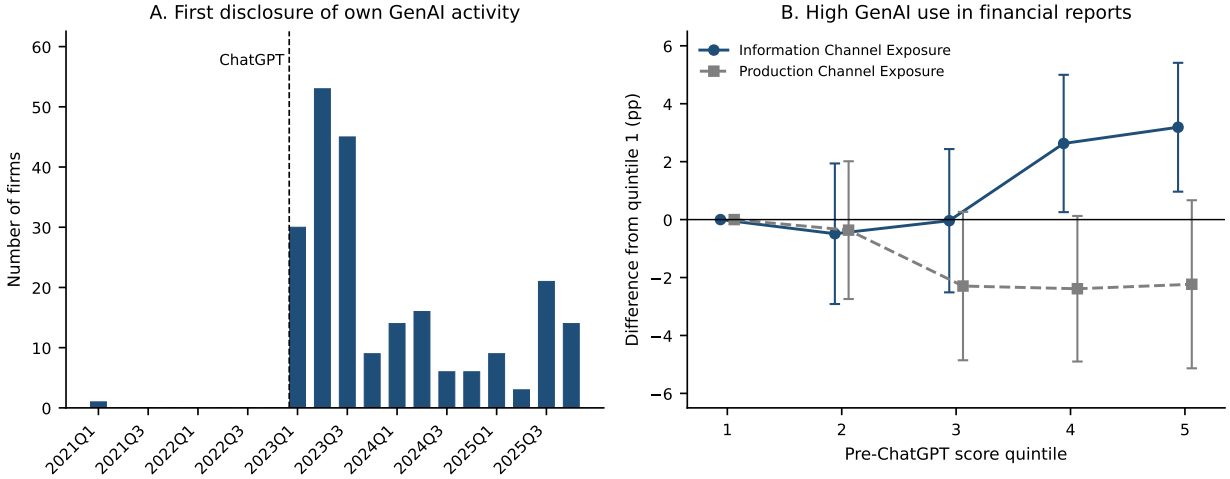
jointly insignificant, and guidance error begins to decline in 2023.[11] Although most annual guidance is issued near the start of the fiscal year, realized revenue depends on decisions made throughout the remainder of that year, so the 2023 decline in guidance error can reflect AI-assisted decisions made after the guidance was issued. Panel B follows the share of new hires in information roles from Table 5, column (1). Before 2023 this share does not move with I_i. It then rises by 0.078 in 2023, 0.072 in 2024, and 0.135 in 2025 per standard deviation of I_i, while production channel exposure leaves it unchanged in every year. In Panel C, sales-per-worker growth likewise changes only after 2022, without a comparable earlier trend, and it moves with both measures. Together, the three panels show that both measures predict higher sales per worker, while only information channel exposure predicts more accurate management forecasts and a shift in hiring toward information roles.

[11]The 2017 coefficient is individually significant. Online Appendix C.1 reports the timing tests for all three panels, and Online Appendix Table A7 reports the placebo date and estimates excluding 2021–2022.

4.5 AI Use after ChatGPT's Release

The guidance, hiring, wage, and sales-per-worker regressions rank firms by what language models could do for them, measured from 2019–2022 earnings calls. We next ask whether greater exposure predicts subsequent use. Figure 2 combines the timing of firms' own GenAI disclosures with an external measure of use in financial reporting. Panel A counts each firm once, in the quarter when it first describes its own GenAI products, operations, workforce, strategy, or investment in an earnings call. Of the 227 first disclosures, 226 occur after ChatGPT became available. These dates record public disclosure, which can differ from the start of operational adoption.

Figure 2: GenAI Disclosures and Use in Financial Reporting



Notes. Panel A reports the quarterly number of firms first describing their own GenAI activity in an earnings call. Each firm enters once, and the passages must have a confidence score of at least 0.70. The dashed line marks the November 2022 release of ChatGPT. Panel B uses the GenAI-authorship probabilities of Blankespoor, deHaan, and Li (2026) for conference calls, risk factors, management discussion and analysis, and earnings press releases. Following their extensive-margin classification, a firm-report pair is classified as high use if the change in the authors' score after ChatGPT's release ever enters the top tercile for that report type and as low use if it always remains in the bottom tercile. Middle-tercile cases are omitted. The high-use measure uses newly written text and excludes unchanged passages. The figure reports coefficients relative to the bottom quintile. Quintiles are formed across all 4,477 scored firms. Regressions include I_i and P_i quintiles jointly, state uncertainty, log assets, two-digit industry effects, and report-type effects. Whiskers give 90 percent confidence intervals from standard errors clustered by firm.

Panel B relates information and production channel exposure to GenAI use measured independently in financial reports by Blankespoor, deHaan, and Li (2026). Their measure estimates the probability that conference calls, risk factors, management discussion and analysis, and earnings press releases were written with generative AI. Together, these documents provide an observable measure of realized firm use in financial reporting rather than the in-

ternal tasks to which a firm applies the technology. Following the authors' extensive-margin classification, a firm-report pair counts as high use when the change in its GenAI-authorship score ever enters the top tercile for that report type. The high-use rate in the bottom information channel exposure quintile is 82.9 percent. Relative to this group, the probability of high use is 3.19 percentage points higher in the top quintile, and the difference is statistically significant. By contrast, no comparable positive relation appears across the production channel exposure quintiles. In the continuous specification, a one-standard-deviation increase in I_i predicts a 1.03 percentage point increase in the probability of high use, whereas the coefficient on P_i is -0.83 percentage points. The two coefficients differ with a p-value of 0.018.

Excluding conference calls leaves the I_i coefficient at 1.16 percentage points and statistically significant, while the P_i coefficient is small and statistically insignificant. When the same relation is estimated over two earlier periods before ChatGPT's release, 2018–2020 and 2020 through November 2022, the I_i coefficient is essentially zero. Online Appendix Table A5 reports these estimates and shows the same pattern with adjusted information channel exposure and with use across multiple report types. Thus, information channel exposure, constructed only from earnings calls held before ChatGPT's release, predicts independently measured firm use afterward. Online Appendix Table A6 further shows that the decline in management guidance error associated with I_i is larger among firms with high use in more report types and has no counterpart under a 2021 placebo break.[12]

The financial-report measure provides a firm-level test of information channel exposure. The Business Trends and Outlook Survey of the U.S. Census Bureau (2026) provides complementary sector-level evidence for production channel exposure. Online Appendix Table A4 aggregates P_i across 18 Census sectors and relates it to the share of businesses reporting AI use in 69 survey waves from 2023 through 2026. Conditional on survey-period effects and workforce exposure, a one-standard-deviation difference in sector-level production channel exposure is associated with a 4.86-percentage-point larger share of businesses reporting AI use, relative to an average share of 9.7 percent. By comparison, the coefficient on workforce exposure is only 0.61 percentage point, and an exact wild-cluster bootstrap across the 18 sectors rejects equality of the two coefficients with a p-value of 0.013. Thus, the two predetermined measures predict subsequent use in different settings: I_i predicts firm-level GenAI

[12]Firms typically announce revenue guidance in earnings press releases and conference calls, two of the four report types scored by Blankespoor, deHaan, and Li (2026).

use in financial reporting, while P_i predicts broader reported AI use across Census sectors.

4.6 Evidence from Analyst Forecasts

Analyst forecasts provide an outside measure of forecastability for the broader I/B/E/S panel. The first column of Table 7 averages analyst disagreement, forecast range, and absolute forecast error, the same three statistics that define opacity. After 2022, this quarterly opacity measure falls by 0.019 standard deviations per standard deviation of I_i. Analyst disagreement and sales-forecast error decline as well, and analyst coverage rises modestly. Production channel exposure, by contrast, does not predict quarterly opacity or management guidance accuracy, although it predicts a decline in analyst sales-forecast error that is not statistically different from the coefficient on I_i. Taken together, the coefficients on I_i suggest that firms with greater information channel exposure became easier to understand from the outside. Management guidance in Table 4 remains the main test of firms' internal information because analysts can also use generative AI tools.

Table 7: Analyst Forecasts after ChatGPT's Release

	Opacity (1)	Disagreement (2)	Range (3)	Sales forecast error (4)	Log analyst coverage (5)
Post $\times I$	-0.019^{**}	-0.019^{**}	-0.015	-0.020^{**}	0.014^{**}
	(0.009)	(0.010)	(0.010)	(0.008)	(0.007)
Post $\times P$	0.013	0.015	0.015	-0.028^{***}	0.013
	(0.012)	(0.012)	(0.013)	(0.011)	(0.009)
p-value: $I = P$	0.060	0.054	0.102	0.574	0.934
Observations	97,513	100,601	111,155	105,551	111,178
Firms	4,051	4,184	4,387	4,201	4,388
Firm FE	Yes	Yes	Yes	Yes	Yes
Two-digit industry × quarter FE	Yes	Yes	Yes	Yes	Yes

Notes. The sample consists of firm-quarter I/B/E/S observations. Columns (1)–(4) use $\log(1+x)$ forecast measures standardized over the pre-2023 sample. Column (1) averages disagreement, range, and absolute error, and column (5) reports log analyst coverage. All regressions include $Post \times I$, $Post \times P$, firm fixed effects, and two-digit industry-by-quarter fixed effects. Standard errors are clustered by firm. Post begins in 2023. $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

5 Estimation and Counterfactuals

In this section, we estimate the assignment model developed in Section 2 using the firm and worker measures introduced in Section 3. The estimated model decomposes the post-2022 rise in sales per worker into the information and production channels, including worker

reassignment. The final subsection allows firms to invest in more signals, which connects the information channel to long-run growth.

5.1 Mapping the Model to Data

We retain the model's two-dimensional output and assignment structure and assign empirical counterparts to firm capabilities, worker skills, decision efficiency, and production efficiency. The firm and worker measures describe the two sides of the labor market, while the forecast and exposure measures determine how ChatGPT changes decision and production efficiency. We allow both changes to differ across firms and estimate the interaction between information and execution tasks. Together, the firm-specific changes and the estimated interaction determine how worker-firm assignment responds.

Firm and worker characteristics – The two assignment dimensions describe firm capabilities and worker skills in information and execution tasks. Because a firm's capability in a task raises the marginal value of worker skill in that task, as in equation (8), we measure each capability by how much of that task the firm's operations involve. Information capability $X^F_{I,i}$ is measured with information fragmentation, which captures how much dispersed material the firm must assemble and compare before deciding. Execution capability $X^F_{E,i}$ is measured with production exposure, which captures the execution tasks in which generative AI tools can reduce labor or time. Both measures are rescaled to lie between zero and one, so zero denotes no measured capability in the corresponding task. On the worker side, information skill is the logarithm of one plus the number of information roles held through 2022, and execution skill is constructed in the same way. Both skills are expressed in units of their sample standard deviations. Among the 6,911 professionals with at least two observed roles, information and execution skills correlate -0.41. By contrast, the two firm capabilities are nearly independent, with a correlation of 0.01.

Information channel – We use analyst-based opacity to set the firm's initial signal precision,

$$p_{A,i} = \exp(\bar{p}_A - \zeta O_i),$$

where $\bar{p}_A$ sets average precision, O_i denotes opacity, and ζ measures how opacity lowers precision. We set prior precision p_0 to one, so expected decision efficiency is $q(p_A) = p_A/(1+p_A)$. We treat the firm-history benchmark, which uses only past sales, as a forecast based on prior information without the additional management signal. Because Lemma 1 makes

mean absolute error proportional to $1-q$, the pre-2023 ratio of management-guidance error to benchmark error sets the mean of $q(p_A)$ to 0.525. Corollary 1 then implies that, holding assignment fixed, more opaque firms have larger output gains from an additional unit of precision.

We next use information exposure to determine how the precision gain varies across firms. With N firms in the estimation sample, firm i's precision increases by

$$\Delta p_{A,i} = \Delta p_A \frac{\exp(\rho_I I_i)}{N^{-1}\sum_j \exp(\rho_I I_j)},$$

where Δp_A is the mean precision increase and ρ_I governs its cross-firm dispersion. As in the regressions, I_i is measured in standard-deviation units.

The output gain from higher precision also depends on the marginal contribution of decision efficiency to output. In equation (3), ℓ_i adds $a_{q,i}$, from the direct decision component, to $a_I X^F_{I,i} X^W_{I,i}$, from the information interaction. Forecast errors identify only this sum, so we measure ℓ_i at the firm's initial match, before any reassignment, using information capability and decision stakes, without restricting the two terms separately:

$$\ell_i = \bar{\ell}\frac{X^F_{I,i}}{\overline{X}^F_I}\exp(\sigma_\ell K_i - \sigma_\ell^2/2),$$

where $\overline{X}^F_I$ is the cross-sectional mean of information capability. We express decision stakes K_i in standard-deviation units, use $\bar{\ell}$ to set the average effect of decision efficiency on output, and let σ_ℓ govern the contribution of decision stakes. Decision efficiency therefore matters more for output when the firm has greater information capability or when its decisions carry greater stakes.

We use mean absolute error in annual management revenue guidance as the empirical counterpart to forecast error. Because annual guidance reflects several decisions made throughout the reporting year, we model annual output as depending on the average efficiency of N_d equally weighted decisions whose errors are independent conditional on the firm. The decision-related component of mean absolute guidance error is then

$$\kappa_d(N_d)\ell_i[1-q(p_{A,i})], \qquad \kappa_d(N_d) = 2\left[F_{\chi^2_{N_d}}(N_d) - F_{\chi^2_{N_d+2}}(N_d)\right],$$

where $F_{\chi^2_k}$ is the distribution function of a chi-squared random variable with k degrees of freedom. Online Appendix Lemma 2 derives this annual-error expression, which reduces to $\kappa_d(1) = 4\varphi(1)$ for one decision. In estimation, N_d may be noninteger and is interpreted as

the effective number of equally weighted independent decisions that yields the same mean absolute annual error. Allowing for forecast error unrelated to the decision, the model-implied mean absolute guidance error in percentage points is

$$\left\{100^2\kappa_d(N_d)^2\ell_i^2[1-q(p_{A,i})]^2+\omega_u^2\right\}^{1/2}.$$

Here ω_u is the mean absolute value of the guidance error unrelated to the decision. The baseline sets $\omega_u = 0$.[13]

Guidance errors identify how precision changes, but their magnitude contains the product $\kappa_d(N_d)\ell_i$. They therefore do not separately determine how much output depends on the decision. Before ChatGPT's release, $\ell_i q(p_{A,i})$ is the part of output attributable to expected decision efficiency. Because the information channel operates through tasks performed by labor, the baseline calibration sets its average as follows:

$$\frac{1}{N}\sum_i \ell_i q(p_{A,i}) = s_L\beta_A,$$

where $s_L = 0.60$ is the labor share and β_A is the share of work tasks a language model can improve. We set $\beta_A = 0.46$, the share of tasks that Eloundou, Manning, Mishkin, and Rock (2024) estimate a language model can improve either directly or through complementary software. In the baseline, this task share sets the average contribution of decision efficiency to output without using the coefficients on information exposure in the sales-per-worker, sales, or employment regressions. The estimated $\Delta p_{A,i}$ then determines the information-channel output gain, and the calibration gives $\bar{\ell} = 0.522$. Since this task share covers both the information and production channels, we later also consider lower values of β_A when we decompose the increase in sales per worker in Section 5.4. As an alternative, we estimate β_A by adding the $Post \times I$ sales-per-worker growth coefficient to the moments. This specification gives $\beta_A = 0.52$, compared with the baseline value of 0.46.

Multiplying a_E, a_I, and c by the same positive number leaves the assignment in Theorem 1 unchanged but changes the term $(X_i^F)'A(q_i, g_i)X_i^W$ in equation (2). For assignment, we therefore set $a_E = 1$, write $w = a_I/a_E$, and express c relative to a_E. For output, we multiply all three by $\bar{\ell}$, the same parameter that sets the average effect of decision efficiency on output,

[13] Combining the two mean absolute errors in squares treats both components as normal. The percentage-point measure uses the first-order correspondence between log changes and percentage changes. Online Appendix Table A20 assigns decision-unrelated forecast errors 25 or 50 percent of pre-2023 forecast-error variance and re-estimates the model. The resulting information-channel shares are 49 and 44 percent.

so that changes in $\bar{\ell}(X_i^F)'A(q_i, g_i)X_i^W$ are proportional changes in output around the initial assignment.

Production channel – We use production exposure P_i in standard-deviation units to determine how much generative AI raises production efficiency. For firm i,

$$g_i = \delta_g(P_i - P_{10})_+,$$

where P_{10} is the tenth percentile of P_i and $(z)_+ = \max\{z, 0\}$. The coefficient δ_g converts the exposure measure into production efficiency, while the floor assigns no measured gain to the least-exposed decile.[14] A higher g_i raises current output both directly and through the complementarity between firm execution capability and worker execution skill. At the initial match, these effects contribute

$$\Delta Q_i^P = g_i + \bar{\ell} g_i X_{E,i}^F X_{E,i}^{W,0}. \tag{23}$$

Here $X_{E,i}^{W,0}$ is the execution skill of the worker assigned to firm i at the initial match. Equation (23) uses worker execution skill in levels and therefore retains its sample mean. The direct term g_i captures the efficiency gain, while the second term captures the additional output from stronger execution complementarity. Unlike a precision gain, a higher g_i leaves management forecast accuracy unchanged. Employment changes by $\eta_L \Delta Q_i^P$, and sales growth equals the sum of sales-per-worker growth and employment growth. Thus η_L measures how employment responds to the production-channel output gain ΔQ_i^P in equation (23). The information channel instead changes workforce skill composition while holding employment fixed.

Assignment – The same changes in q_i and g_i that raise output also change the worker a firm hires. Equation (9) gives firm i's marginal values of worker skills as $m_i^F = A(q_i, g_i)'X_i^F$. Under the normalization $a_E = 1$, these values before and after ChatGPT's release are

$$\begin{aligned} m_i^{F,0} &= \left(X_{E,i}^F + cX_{I,i}^F,\ wq(p_{A,i})X_{I,i}^F + cX_{E,i}^F\right)', \\ m_i^{F,1} &= \left((1+g_i)X_{E,i}^F + cX_{I,i}^F,\ wq(p_{A,i}+\Delta p_{A,i})X_{I,i}^F + cX_{E,i}^F\right)'. \end{aligned}$$

To represent firms of different sizes, we treat each observed firm as a firm type with mass proportional to its mean employment from 2019 through 2022. Each firm type is matched with an equal mass of workers, preserving the one-to-one assignment structure in Section 2.

[14] Online Appendix B.2 derives g_i from the share of the firm's production operations that generative AI tools can perform.

Theorem 1 assumes Gaussian marginal values, but firm-specific $p_{A,i}$ and g_i generally produce a non-Gaussian distribution, which we approximate by a Gaussian with the same employment-weighted covariance Σ_m. Applying the map in Theorem 1 to this approximation and the worker distribution with observed covariance Σ_W then determines how workers are distributed across firm types under $m_i^{F,0}$ and under $m_i^{F,1}$.

With both assignments in hand, we separate the output change holding assignment fixed from the contribution of worker reassignment, the two terms in equation (16). For reassignment, writing marginal values as their mean plus a firm-specific deviation separates the value of worker skill that is common across firms from the value specific to a match. The common part cannot affect assignment, and its aggregate contribution is unchanged because the assignment preserves the worker distribution. We therefore value reassignment with marginal values relative to their employment-weighted mean and use the same firm masses when aggregating output, as in the decomposition in Online Appendix C.2.

5.2 Wage Income

The H-1B wage response of information hires to I_i measures how their pay changes from 2023 onward. We obtain its model counterpart from the equilibrium wage schedule.

Wage schedule – The proof of Theorem 1 constructs the quadratic wage schedule that supports the assignment. With a_E, a_I, and c multiplied by $\bar{\ell}$, the wage schedule is

$$W(X^W) = \bar{W} + \bar{\ell}\left[\mu_m' x^W + \frac{1}{2}(x^W)' B^{-1} x^W\right].$$

Its gradient is $\bar{\ell}$ times the marginal-value vector of the firm matched with that worker. The linear term values worker skills at the mean marginal value μ_m. The quadratic term adds the matched firm's deviation from that mean, $m^F - \mu_m = B^{-1}x^W$, and therefore depends on worker-firm assignment.

Wage changes – The information component of the gradient of W is the marginal return to worker information skill,

$$\frac{\partial W(X^W)}{\partial X_I^W} = \bar{\ell}\left\{\mu_{m,I} + \left[B^{-1}x^W\right]_I\right\}.$$

It rises when firms that place greater value on information skill are matched with more information-skilled workers. Because the H-1B wage regression compares information hires within occupation, state, and year, its model counterpart is the change in the quadratic

term of the wage schedule. We convert that change into log wages by dividing by the labor share.[15]

5.3 Estimation and Model Fit

Parameters – We set $\sigma_\ell = 0.30$, so a one-standard-deviation increase in decision stakes raises ℓ_i by 35 percent.[16] We estimate the remaining eight parameters,

$$\psi = (N_d, \zeta, \Delta p_A, w, \delta_g, \eta_L, c, \rho_I)'.$$

Estimation – To construct the model counterparts of the empirical moments, we regress model outcomes on the same I_i and P_i used in the data. We estimate ψ by minimum distance using ten moments in three groups.[17] First, guidance errors depend on decision precision, so six guidance moments identify the information parameters: the pre-2023 firm-level error intercept and its opacity slope, the post-2022 coefficients on I_i, the F-K-I composite, and opacity, and the coefficient on I_i for the shrinkage-benchmark error. Second, two labor moments discipline worker assignment. Because Capital IQ professional histories do not observe every worker-firm match, we summarize initial sorting with the relation between information fragmentation and the share of new hires in information roles before ChatGPT's release. The H-1B wage response of information hires to I_i measures the change in the marginal return to information skill. Third, firm performance identifies the production parameters through the $Post \times P$ coefficients for sales per worker and employment, both measured as annual differences over 2023–2025.[18]

The sign of the cross-task interaction – We use the H-1B wages of information and execution hires to determine whether firm capability in one task raises or lowers the marginal

[15] At the initial assignment, labor income is s_L of output. A first-order change in the wage W, measured as a share of output, therefore maps into a log wage change of $\Delta W/s_L$.

[16] Setting $\sigma_\ell = 0$ lowers the information-channel share from 49.2 to 45.6 percent in Online Appendix Table A20.

[17] The error intercept and opacity slope come from regressing each firm's mean pre-2023 guidance error on opacity, with one observation per firm. The intercept is 9.827, while the observation-weighted mean for the same years in Table 4 is 7.864. The management-guidance-error coefficients on the F-K-I composite and opacity, -1.485 and -0.357, come from a joint regression that also includes $Post \times P$. The slope of the pre-2023 share of Capital IQ new hires in information roles on information fragmentation is -0.005, controlling for production channel exposure P_i. Online Appendix C.2 describes the weighting matrix and the 300 joint draws of the empirical moments that give the standard errors and prediction intervals.

[18] Each coefficient therefore corresponds to one third of the model-implied change in sales per worker or employment over 2023–2025. The counterfactuals in Section 5.4 later report the full three-year change in sales per worker, averaged across firms with weights proportional to their mean employment from 2019 through 2022.

value of worker skill in the other. Panel C of Online Appendix Table A13 shows that offered wages for information hires rise by 0.021 log units more per standard deviation of I_i than per standard deviation of P_i, with a p-value of 0.045 for equality. Production exposure thus does not raise the return to information skill, unlike information exposure. Because c also links firm information capability with worker execution skill, a separate specification adds three moments and leaves c unrestricted. These moments are the coefficient on P_i for information-hire wages and the coefficients on I_i and P_i for execution-hire wages. This specification gives $c = -1.01$, so the baseline specification, which uses the ten moments above, imposes $c \leq 0$.[19]

Non-targeted moments – We leave two sets of moments out of estimation: the $Post \times I$ coefficient for the share of Capital IQ new hires in information roles, which tests the model's prediction for worker reassignment, and the $Post \times I$ coefficients for sales per worker, sales, and employment, which test its prediction for firm performance.

Model fit – The Jacobian has full column rank, establishing local identification, and the overidentifying restrictions do not reject the model: $J = 3.55$ on two degrees of freedom with a p-value of 0.17. Turning to the estimates in Table 8, the estimated precision change raises average expected decision efficiency by 0.061 from a pre-2023 mean of 0.525. The estimate $N_d = 7.31$ maps annual guidance into the same decision uncertainty as about seven equally weighted independent decisions. Finally, the baseline estimate $c = -1.274$ means that a firm's capability in one task lowers its marginal value of worker skill in the other task, and this interaction matters for the fit: setting c to zero raises the minimum-distance criterion from 3.55 to 230.5.[20] Online Appendix Table A18 reports 95 percent profile ranges for c and four other parameters.

Panel C of Table 8 compares the four non-targeted moments with the model predictions. For worker reassignment, the model predicts an increase of 0.108 in the share of new hires in information roles, compared with 0.105 in the data. Under the baseline calibration, it predicts 2.110 percentage points faster annual sales-per-worker growth, compared with 2.074 in the data. Because the information channel reallocates a fixed workforce across tasks, the model predicts little change in total employment. For sales and employment growth,

[19] Online Appendix Table A20 shows that the unrestricted specification gives an information-channel contribution of 5.6 percentage points and a 44.4 percent share of the cumulative increase in sales per worker, compared with 6.0 percentage points and 49.2 percent in the baseline.

[20] The Gaussian assignment requires only a nonsingular Σ_m, so $A(q_i, g_i)$ can be indefinite, as it is at the sample mean under the estimated c. Sorting is then positive in the marginal-value vector $m_i^F = A(q_i, g_i)' X_i^F$, as in equation (10), although it need not be positive within each task.

the model predicts increases of 2.134 and 0.024 percentage points, compared with 2.053 and −0.021 in the data. Together, the hiring, sales-per-worker growth, and employment predictions are consistent with the data, with a joint p-value of 0.997.[21] Re-estimating the model after omitting each targeted moment in turn gives sales-per-worker growth predictions between 1.42 and 2.46 in Online Appendix Table A17.

The fitted model implies that management-guidance error falls by 2.06 percentage points on average. In the data, among the 699 firms that issued guidance both before 2023 and afterward, raw mean absolute error falls from 8.18 to 6.29 percentage points, a decline of 1.89 points. Panel C of Table 8 reports the two magnitudes side by side, and their similarity provides a separate check on how the model translates signal precision into guidance error.

The model also closely matches sales-per-worker growth among firms with the greatest information channel exposure. Online Appendix Table A19 compares firms in the top quintile of information channel exposure with the lower four and evaluates the model on the same firms. Annual sales-per-worker growth is 4.86 percentage points higher for the top quintile in the data and 4.97 points higher in the model. The point estimates are close, and the equality test has a p-value of 0.95. For the share of new hires in information roles, however, the model predicts a larger increase, 32.4 percentage points compared with 23.7 points in the data, with a p-value of 0.15 for equality.

5.4 Counterfactuals

We use the estimated model for two counterfactuals. The first measures the cross-firm difference in sales per worker associated with information channel exposure, holding production channel exposure at its mean. The second measures the cumulative sample-wide increase in sales per worker from 2023 to 2025 relative to keeping signal precision and production efficiency at their values before ChatGPT's release and decomposes that increase across the two channels.

Panel A of Table 9 shows that moving from the 10th to the 90th percentile of I raises predicted sales per worker by 13.6 percentage points. Of this difference, 9.1 percentage points arise with assignment held fixed and 4.5 from worker reassignment. Thus the same cross-firm variation that predicts improvements in guidance changes both decision efficiency and the allocation of information skill.

[21] The joint test omits the sales coefficient because sales-per-worker growth equals sales growth minus employment growth.

Table 8: Model Estimates and Non-Targeted Moments

Panel A. Estimates and implied decision-efficiency gain			
Quantity	Estimate	Quantity	Estimate
N_d, decisions per guidance year	7.311	ζ, opacity-to-precision slope	0.388
	(0.819)		(0.059)
Implied mean gain in decision efficiency	0.061	ρ_I, precision-gain heterogeneity	1.894
	(0.023)		(0.862)
w, relative information weight	1.587	c, cross-task interaction	−1.274
	(0.338)		(0.192)
δ_g, production-efficiency coefficient	0.028	η_L, employment response	−0.654
	(0.012)		(0.672)

Panel B. Targeted moments			
Moment	Data	Model	Std. resid.
Pre-2023 guidance-error intercept (pp)	9.827	9.708	−0.28
	(0.434)		
Guidance-error–opacity slope	2.672	2.439	−0.55
	(0.425)		
Management-guidance error: Post $\times I$	−1.617	−1.632	−0.06
	(0.271)		
Management-guidance error: Post $\times$ F-K-I composite	−1.485	−1.386	0.43
	(0.231)		
Management-guidance error: Post $\times$ opacity	−0.357	−0.398	−0.14
	(0.301)		
Shrinkage error: Post $\times I$	−0.234	0.000	0.71
	(0.332)		
Information-hire wage: Post $\times I$	0.014	0.011	−0.73
	(0.005)		
Sales/worker: Post $\times P$	1.528	1.638	0.15
	(0.721)		
Employment: Post $\times P$	−0.992	−0.994	0.00
	(0.477)		
Pre-2023 slope: share of new hires in information roles	−0.005	−0.001	0.36
	(0.009)		

Panel C. Non-targeted moments			
Moment	Data	Model [90% interval]	Std. resid.
Share of new hires in information roles: Post $\times I$	0.105	0.108	0.20
	(0.013)	[0.081, 0.129]	
Sales/worker: Post $\times I$	2.074	2.110	0.05
	(0.501)	[1.439, 2.838]	
Sales: Post $\times I$	2.053	2.134	0.11
	(0.547)	[1.467, 2.857]	
Employment: Post $\times I$	−0.021	0.024	0.13
	(0.355)	[0.005, 0.043]	
Descriptive guidance-error change			
Paired-sample guidance-error change	−1.890	−2.063	–
	(0.380)		

Panel D. Specification			
Targeted moments	10	Parameters	8
Jacobian rank	8	J statistic (2 d.f.)	3.55
Overidentification p-value	0.169	Sales/worker growth range	1.42–2.46
Joint non-targeted test, p-value	0.997		

Notes. The F-K-I composite is the geometric mean of information fragmentation, decision stakes, and information channel exposure, expressed in standard-deviation units. Parentheses contain standard errors and brackets contain 90% prediction intervals, based on 300 joint draws of the empirical moments. Standardized residuals divide the difference between the model and the data by the data standard error in Panel B and by the standard deviation of that difference in Panel C. Baseline decision efficiency treats the firm-history benchmark as a forecast based on prior information and uses the pre-2023 ratio of management-guidance error to benchmark error. The baseline calibration uses a labor share of 0.60 and the 0.46 task share in Eloundou, Manning, Mishkin, and Rock (2024), which includes tasks improved directly or through complementary software. Growth moments are in percentage points, and the information channel holds employment fixed. The Panel C joint test uses the hiring-share, sales-per-worker, and employment moments. The sales-per-worker growth range omits each targeted moment in turn.

Panel B shows that, under the baseline calibration, the estimated increases in signal precision and production efficiency raise predicted sales per worker in the sample by a cumulative 12.3 percentage points from 2023 to 2025. The information channel contributes 6.0 percentage points, accounting for 49.2 percent of this increase. Of this information-channel contribution, 4.1 percentage points arise with assignment held fixed and 1.9 percentage points arise from reassignment. Under the baseline Gaussian assignment, reassignment therefore accounts for 31 percent of the information-channel contribution. By comparison, the production channel contributes 6.1 percentage points, and the interaction between the channels through assignment adds 0.1 percentage point.

Task-share calibration – Panel C shows the sensitivity to alternative task-share calibrations. The lower calibrations predict annual sales-per-worker growth coefficients of 0.94 to 1.42 percentage points and assign 28 to 38 percent of the cumulative increase to the information channel. Estimating the task share with the empirical $Post \times I$ sales-per-worker growth coefficient as an additional moment gives an information-channel share of 52.0 percent. Thus the precise share depends on β_A, but each calibration attributes a substantial part of the cumulative increase in sales per worker to the information channel.

The information-channel share is more sensitive to initial precision and the production channel exposure floor than to the assignment calculation. Online Appendix Table A20 gives a range of 35 to 64 percent under alternative values of initial precision and the exposure floor. By comparison, solving the discrete assignment problem on the observed firm sample lowers the reassignment contribution from 1.9 to 1.7 percentage points and the information-channel share only to 48.6 percent. Requiring the interaction matrix at the sample mean to be positive definite lowers the information-channel share to 46.0 percent and yields a small reassignment contribution. However, the overidentifying restrictions reject this specification, with $J = 16.91$ and a p-value below 0.001. Requiring every firm's initial interaction matrix to be positive definite lowers the information-channel share further, to 40.2 percent. The reassignment contribution is therefore less stable than the overall information-channel share.

The counterfactuals in Table 9 treat the ChatGPT release as a one-time increase in signal precision and production efficiency. Online Appendix Table A16 follows the fitted 2025 allocation forward when both increases remain fixed and firm capabilities continue to evolve according to equation (4). When learning from matched workers accounts for 30 percent of the 2014–2022 year-to-year persistence in firms' sales per worker, as in the baseline,

the cumulative increase in sales per worker rises from 12.3 percentage points through 2025 to 17.0 percentage points in the long run. Shares of 15 and 45 percent give increases of 13.7 and 18.8 percentage points.

Table 9: Information, Assignment, and Production Counterfactuals

Panel A. Information channel exposure across firms

	Total difference	Assignment held fixed	Reassignment
90th versus 10th percentile of I	13.6	9.1	4.5

Panel B. Cumulative 2023–2025 level changes

	Sales/worker (pp)	Guidance error change (pp)	Information skill slope	Share of combined (%)
Combined	12.3	−2.06	0.39	100.0
	(3.1)			
Information channel	6.0	−2.06	0.38	49.2
	(1.2)			
Assignment held fixed	4.1	−2.06	0.00	33.8
	(0.9)			
Reassignment	1.9	0.00	0.38	15.4
	(0.6)			
Production channel	6.1	0.00	0.00	50.0
	(2.6)			
Interaction	0.1	0.00	0.01	0.8

Panel C. Alternative task-share calibrations

β_A	$\bar{\ell}$	Sales/worker growth	Combined increase	Information channel share (%)	Reassignment share (%)
0.20	0.227	0.936	9.5	28.0	33.2
0.30	0.340	1.421	10.6	37.7	33.5
0.46	0.522	2.110	12.3	49.2	31.3
0.52	0.593	2.204	12.3	52.0	29.3
(estimated)					
Data		2.074			
		(0.501)			

Notes. Panel A holds production channel exposure at its sample mean and reports model-implied sales-per-worker differences in percentage points. Panel B reports cumulative 2023–2025 differences in levels relative to no increase in signal precision or production efficiency. Sales-per-worker entries are percentage-point differences for the sample rather than economy-wide productivity growth. The information-skill slope is the change in the assigned worker's information skill per standard deviation of I, conditional on P. Parentheses contain standard errors. Reassignment is the information-channel contribution minus its contribution with assignment held fixed. The interaction row is the combined increase less the two separate channel contributions. The production-channel contribution includes the direct efficiency gain and the change in the execution interaction in equation (23). Panel C re-estimates the model at alternative task shares. The 0.30 and 0.46 values correspond to alternative exposure definitions in Eloundou, Manning, Mishkin, and Rock (2024). The 0.20 row provides a lower calibration, and 0.46 is the baseline calibration. The estimated row uses the $Post \times I$ sales-per-worker coefficient as an additional moment. Components may not sum exactly because of rounding.

5.5 Information Investment and Long-Run Growth

In the estimated model, a one-time increase in signal precision has only a transitory effect on growth. Output rises as worker assignment and firm capabilities adjust, but eventually converges to a higher level. To account for long-run growth, we follow Jovanovic and Ma (2023) and allow firms to invest in additional signals. This investment makes the direct decision component $a_{q,i}q_i$ in equation (2) endogenous while preserving the same signal-extraction and assignment problems.

Information investment – After forming a match and before observing its signals, firm i invests $n_{i,t} \geq 0$ units of output in additional signals. Each signal informs a separate decision about a project, product, or operating choice, and each such decision adds q_i to expected output. The coefficient $a_{q,i}$, which in equation (2) measures how much output depends directly on decision efficiency, therefore becomes the number of signals the firm holds. Starting from the number in the preceding period, the number available at date t is

$$a_{q,i,t} = \phi_a a_{q,i,t-1} + \chi n_{i,t}^{\alpha} a_{q,i,t-1}^{1-\alpha}, \qquad 0 < \alpha < 1, \tag{24}$$

where $\phi_a > 0$ governs how many of last period's signals carry over, $\chi > 0$ measures the productivity of investment, and α is the elasticity of new signals with respect to investment. These signals enter current output through $q_i a_{q,i,t}$. Because this term does not depend on the worker's skill, the investment choice does not alter the assignment characterized in Theorem 1. It instead changes how many decisions benefit from the information available at that match.

The firm chooses investment to maximize the direct decision component of current output net of its cost,

$$\max_{n_{i,t} \geq 0} \left\{ q_i a_{q,i,t} - n_{i,t} \right\}. \tag{25}$$

The investment problem retains the bounded decision-efficiency function in equation (1). By raising q_i, AI raises the return to each additional signal.

Proposition 4 (Information investment and growth) *Suppose firm i invests in signals according to equation (25). Its investment relative to the number of signals in the preceding period is*

$$\iota_i \equiv \frac{n_{i,t}}{a_{q,i,t-1}} = (\alpha \chi q_i)^{1/(1-\alpha)}. \tag{26}$$

The number of signals then grows by the factor

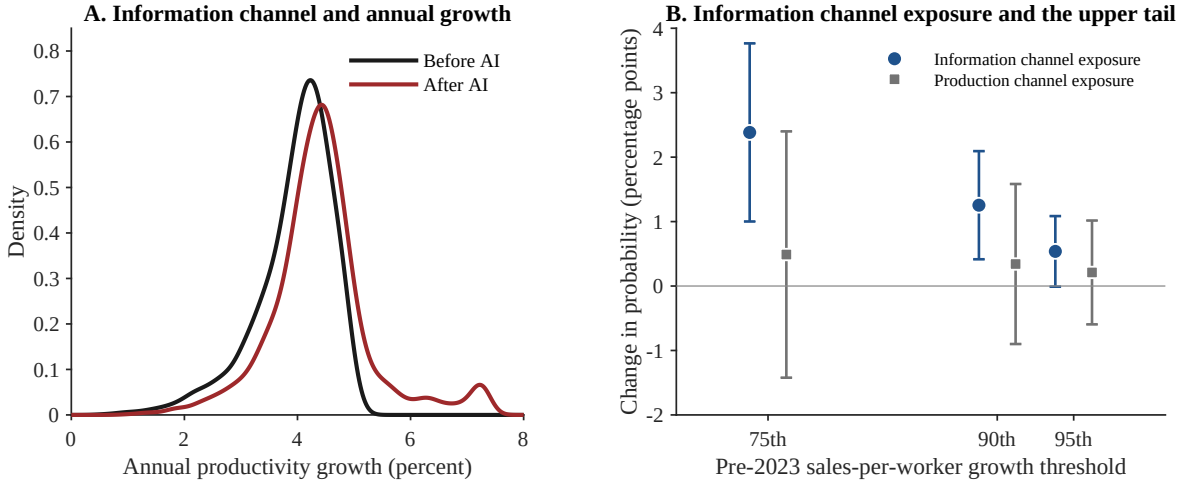
$$G_i \equiv \frac{a_{q,i,t}}{a_{q,i,t-1}} = \phi_a + \chi \left(\alpha \chi q_i\right)^{\alpha/(1-\alpha)}. \tag{27}$$

Both ι_i and G_i increase with expected decision efficiency. Thus more precise information raises investment in signals and long-run growth. Once the direct decision component dominates the other terms of equation (2), which remain stationary, output per worker grows at rate $G_i - 1$ per period.

Proof. See Online Appendix B.

Proposition 4 turns the information channel into an annual growth mechanism. More precise signals improve current decisions and also raise the return to additional signals. Investment then increases the number of decisions that signals inform.

Figure 3: Information Investment and Annual Firm Growth



Notes. Panel A plots densities of model-implied annual productivity growth at initial precision and after the estimated precision increase for 4,477 firms. We set $\alpha = 0.5$ and $\phi_a = 1$, and choose χ so that median model-implied annual growth at initial precision is 4.1 percent. The fitted initial signal precision and its estimated increase determine q_i in each period. Panel B reports coefficients and 95 percent confidence intervals on $Post \times I_i$ and $Post \times P_i$. The dependent variable indicates whether sales-per-worker growth exceeds the stated percentile of its pre-2023 distribution. The coefficients correspond to a one-standard-deviation increase in each exposure. Regressions use 22,236 observations for 3,109 firms, remove each firm's projected 2018–2022 trend, and include firm and two-digit industry-by-year effects. Standard errors are clustered by firm.

Quantitative implications – We use the fitted precision of each firm's signal and its estimated increase to evaluate equation (27), treating each period as one year. We set $\alpha = 0.5$ and $\phi_a = 1$, then choose χ so that median model-implied annual growth at initial precision

is 4.1 percent, equal to median annual sales-per-worker growth from 2018 through 2022.[22] At this value of χ, the median investment cost at initial precision is 2.0 percent of the direct decision component of output and a smaller share of total output, which also includes the match and execution components. The information channel alone raises annual growth by a median of 0.18 percentage point. Because the precision gain rises with information channel exposure, the increase in annual growth is concentrated among the most exposed firms and exceeds 1.30 percentage points for one firm in ten. These increases expand the upper tail of the growth distribution in Panel A of Figure 3, raising the share of firms whose predicted growth exceeds the 90th percentile of growth at initial precision from 10.0 to 28.5 percent.[23]

Panel B of Figure 3 asks whether the upper tail already expands in the data. A one-standard-deviation increase in I_i is associated with a 2.38-percentage-point higher probability of exceeding the pre-2023 75th percentile of sales-per-worker growth, from a base of 25 percent. The probability of exceeding the 90th percentile also rises, by 1.25 percentage points from 10 percent. The corresponding production channel exposure coefficients are 0.49 and 0.34 percentage points and are estimated less precisely. The difference between the two exposure coefficients is also imprecise. Online Appendix Table A15 reports additional thresholds and the amount by which growth exceeds each threshold.[24]

6 Robustness and Discussion

Since ChatGPT's release, firms with greater information channel exposure issue more accurate revenue guidance, shift hiring toward information roles, and raise sales per worker. We first ask whether conditions around November 2022, industry composition, or firm characteristics such as uncertainty and prior technology adoption could produce these changes without any role for language models. We then examine whether the hiring and wage estimates depend on how the worker data are treated. Finally, we connect the guidance estimates to decision efficiency.

[22] This calibration treats observed growth over these years as the long-run rate $G_i - 1$.

[23] With $\alpha = 0.5$ and $\phi_a = 1$, equation (27) gives $G_i - 1 = \chi^2 q_i/2$, so each firm's growth rate is proportional to q_i. The increases in annual growth therefore scale in proportion to the 4.1 percent median growth rate used to set χ, whereas the share of firms above that 90th percentile does not depend on that rate.

[24] Our analysis focuses on the upper tail generated by the information channel. Gans and Landry (2026) instead show how worker incentives can slow investment in AI training and automation.

6.1 Timing, Macroeconomic Controls, and Industry Composition

Firms with high I_i could have been hit harder by the pandemic, over-hired in 2021 and 2022, or made unusually large forecast errors just before ChatGPT became available, so part of the post-2022 change could be mean reversion. We use the alternative specifications in Online Appendix Table A7 to test whether these earlier differences explain the coefficients on I_i from 2023 onward.

Pandemic exposure and forecast errors – Interacting Post with pandemic-era sales volatility, opacity, and the firm's 2021–22 guidance error leaves the guidance coefficient on I_i at -1.676, close to the baseline estimate of -1.617.

Mean reversion – Firms whose guidance was less accurate before 2023 have more room to improve. To address this concern, we compare firm means before and after 2022 in Online Appendix Table A10 and condition each firm's later error on its earlier error. The coefficient on I_i is -1.618 percentage points, nearly identical to the baseline estimate. Across nonlinear controls for the earlier error, trimming, and a requirement of at least three pre-2023 years of guidance, the coefficient ranges from -1.320 to -2.006 percentage points. By contrast, in a placebo comparing 2020–2022 with 2017–2019, the corresponding coefficient is 0.057 and not statistically different from zero.

Over-hiring – Interacting Post with pandemic-era volatility, 2020–22 over-hiring, and 2019–22 growth in sales per worker leaves the sales-per-worker coefficient on I_i at 1.522, about one-quarter below the baseline estimate of 2.074 but still economically sizable.

Removing 2021–22 – Dropping these years entirely gives a coefficient on I_i of -1.396 in the management-guidance regression. The corresponding sales-per-worker specification uses unadjusted growth and yields 1.873, close to the unadjusted estimate of 1.922.

Placebo dates – Moving the break to 2019 for guidance and to 2016 for sales per worker gives 0.207 and -0.131, neither distinguishable from zero.

Federal Reserve tightening – Because ChatGPT's release followed the Federal Reserve's rapid rate increases in 2022, we also estimate the association between I_i and guidance error from April 1 through November 29, 2022. Panel C of Online Appendix Table A7 shows that the coefficient for this window is 0.359 and statistically indistinguishable from zero, compared with -1.596 from 2023 onward.

Taken together, the robustness tests and the joint pretrend tests in Figure 1 place the break at the end of 2022, when ChatGPT became available, not in the recovery from the

pandemic or the 2022 tightening cycle. The timing of firms' first GenAI disclosures in Figure 2 reinforces this conclusion.

A related concern is that firms with high I_i are concentrated in sectors that benefited after 2022. For sales per worker, absorbing changes within three-digit industries lowers the coefficient on I_i from 2.074 to 1.508, but it remains statistically significant. With four-digit industry-by-year effects, the coefficient is 0.949 and significant at the 10 percent level, compared with 2.079 on the same sample, as Online Appendix Table A8 reports.

For guidance, we assign each firm its modal Compustat NAICS code across years, which gives a common classification at two, three, and four digits. This classification is available for a smaller sample, within which the guidance coefficient in Panel A of Online Appendix Table A8 instead becomes more negative as the industry effects become more detailed, from -2.266 with two-digit industry-by-quarter effects to -2.702 with four-digit effects on a common sample of 502 firms.

Panel C of Online Appendix Table A8 further provides a direct within-industry test by permuting I_i and P_i jointly across firms within each four-digit industry. None of 400 permutations produces a guidance coefficient as negative as the actual coefficient of -2.702, giving a p-value below 0.0025. Differences across industries therefore do not reproduce the improvement in management guidance among firms with greater I_i.

The guidance and sales-per-worker coefficients on I_i are also similar after excluding sectors closely tied to digital technology. Following Babina, Fedyk, He, and Hodson (2024), we exclude information and professional services and obtain coefficients of -1.850 and 2.319. A narrower exclusion of data processing, publishing, accounting, and travel-reservation industries gives sales-per-worker coefficients of 2.197 for I_i and 1.821 for P_i. Online Appendix Tables A13 and A14 report these estimates and the corresponding results for margins, operating expenses relative to sales, and operating income per worker.

6.2 Information Channel Exposure and Other Information Characteristics

To separate language-model fit from the other information characteristics, column (5) of Online Appendix Table A7 enters I_i jointly with the U-F-K composite, the geometric mean of uncertainty, fragmentation, and stakes expressed in standard-deviation units. Information channel exposure still predicts a 1.617 percentage-point decline in guidance error and a 2.072

percentage-point rise in sales-per-worker growth, whereas the corresponding coefficients on the composite are −0.007 for guidance error and −0.097 for sales-per-worker growth. The near-zero coefficients on the composite distinguish language-model fit from the other characteristics of the firm's information problem, consistent with the predetermined validation in Table 3.

The baseline specification enters I_i and P_i jointly and does not fold uncertainty, fragmentation, or stakes into I_i. Leaving these characteristics outside I_i keeps information channel exposure focused on language-model fit. As a complementary specification, Online Appendix Table A11 replaces I_i with adjusted information channel exposure, I_i^{adj}, which also incorporates uncertainty, fragmentation, and stakes. The adjusted information channel exposure measure predicts lower forecast errors, more information hiring, higher information-hire wages, and faster sales-per-worker growth.

We next ask whether I_i is distinct from conventional measures of analyst disagreement. In Online Appendix Table A12, adding disagreement leaves the coefficients on I_i at −1.595, 0.104, and 2.119 for guidance error, the share of new hires in information roles, and sales-per-worker growth, while disagreement itself is statistically insignificant in all three regressions. We then interact information channel exposure with the broader analyst-based opacity measure. Greater opacity amplifies both the decline in guidance error and the increase in sales-per-worker growth, although its interaction with the share of new hires in information roles is small and imprecise.

Adding measures of prior technology adoption leaves the guidance-error, hiring, and sales-per-worker growth coefficients nearly unchanged. As Panel B of Table 6 shows, the sales-per-worker coefficient on I_i remains 1.973 when prior AI hiring and its interaction with opacity enter jointly. Entering industry AI exposure, prior AI hiring, and R&D intensity together leaves the information channel exposure coefficients at −2.037 in the guidance regression and 1.841 in the sales-per-worker regression, as Online Appendix Table A8 reports. The workforce exposure measure of Eisfeldt, Schubert, Zhang, and Taska (2026) leads to the same conclusion. On the firms it covers, adding its interaction with Post changes the information channel exposure coefficients only from −1.344 to −1.301 for guidance error, from 0.118 to 0.120 for the share of new hires in information roles, and from 2.145 to 2.138 for sales-per-worker growth, as Panel B of Online Appendix Table A3 shows.

6.3 Worker Measurement

A new hire is each person's first observed role at the firm, so promotions and other internal role changes do not enter the share of new hires in information roles or hire counts. Online Appendix Table A13 shows that the hiring and wage estimates are similar under alternative treatments of the worker data. Weighting observations by classified hires and adding two-digit industry-by-year effects gives a coefficient of 0.109 on the same firm-years as the baseline hiring-share regression, and ending the sample in 2023 or 2024 gives 0.093 and 0.089. The ten largest H-1B filers account for 49.4 percent of applications for information roles, but dropping them changes the information-hire wage coefficient only from 0.012 to 0.014 after absorbing occupation-by-year effects.

6.4 From Guidance Error to Decision Efficiency

The signal in the model concerns the payoff-relevant state behind a firm's information-intensive decision, whereas the empirical measure is the error in management's own revenue guidance. Lemma 1 links the two because the decision-related component of mean absolute forecast error is proportional to $1 - q(p_A)$. This component is scaled by ℓ_i, the marginal contribution of decision efficiency to output in equation (3), so an increase in precision that lowers guidance error also raises output. The model therefore attributes the decline in guidance error in Table 4 and the faster growth in sales and sales per worker in Table 6 among firms with greater information channel exposure to the same increase in precision.

7 Conclusion

This paper documents an information-based source of AI-led growth. Firms with greater predetermined information channel exposure improve their forecasts, increase information hiring relative to execution hiring, and raise sales without adding workers. Information exposure also predicts later GenAI use in independently scored reports, including reports other than earnings calls. Production channel exposure, by contrast, predicts higher sales per worker largely through lower employment. Without targeting either outcome, the estimated assignment model predicts the shift toward information hiring and the faster sales-per-worker growth among firms with greater information channel exposure. Under the baseline calibration, the information channel accounts for about half of the post-2022 rise in sales per worker.

The evidence on guidance, labor allocation, sales, and employment, together with the model decomposition, distinguishes gains from better decisions and worker-firm assignment from gains in production efficiency. When firms invest in more signals, greater signal precision also raises long-run growth and expands the upper tail of the distribution of firm growth.

References

[1] Acemoglu, D. (2025). The simple macroeconomics of AI. *Economic Policy*, 40(121), 13–58.

[2] Albagli, E., Hellwig, C., & Tsyvinski, A. (2023). Imperfect financial markets and investment inefficiencies. *American Economic Review*, 113(9), 2323–2354.

[3] Altonji, J. G., & Pierret, C. R. (2001). Employer learning and statistical discrimination. *Quarterly Journal of Economics*, 116(1), 313–350.

[4] Ash, E., Hansen, S., & Muvdi, Y. (2024). Large language models in economics (cepr Discussion Paper 19479).

[5] Autor, D., Chin, C., Salomons, A. M., & Seegmiller, B. (2026). What makes new work different from more work? (nber w34986).

[6] Babina, T., Fedyk, A., He, A., & Hodson, J. (2024). Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics*, 151, 103745.

[7] Bali, T. G., Kelly, B. T., Moerke, M., & Rahman, J. (2023). Machine forecast disagreement (nber w31583).

[8] Barron, O. E., Kim, O., Lim, S. C., & Stevens, D. E. (1998). Using analysts' forecasts to measure properties of analysts' information environment. *The Accounting Review*, 73(4), 421–433.

[9] Benmelech, E., & Zator, M. (2022). Robots and firm investment (nber w29676).

[10] Bhandari, A., Jugulum, S., Veldkamp, L., & Venkateswaran, V. (2026). Firm specialization vs. diversification (Working paper).

[11] Blankespoor, E., deHaan, E., & Li, Q. (2026). Generative AI in financial reporting. *Journal of Accounting Research*, 64(3), 1189–1232.

[12] Bonelli, M. (2026). Data-driven investors. *Review of Financial Studies*, 39(7), 1909–1969.

[13] Brynjolfsson, E., Li, D., & Raymond, L. R. (2023). Generative AI at work (nber w31161).

[14] Bybee, L. (2024). The ghost in the machine: Generating beliefs with large language models (Working paper).

[15] Carlson, J., & Dell, M. (2025). A unifying framework for robust and efficient inference with unstructured data (Working paper).

[16] Crane, L. D., Karra, A., & Soto, P. E. (2025). Total recall? Evaluating the macroeconomic knowledge of large language models (feds 2025-044).

[17] Czarnitzki, D., Fernandez, G. P., & Rammer, C. (2022). Artificial intelligence and firm-level productivity (zew Discussion Paper 22-005).

[18] Diether, K. B., Malloy, C. J., & Scherbina, A. (2002). Differences of opinion and the cross section of stock returns. *Journal of Finance*, 57(5), 2113–2141.

[19] Edmond, C., & Mongey, S. (2026). Unbundling labor (nber w35783).

[20] Eisfeldt, A. L., Schubert, G., Zhang, M. B., & Taska, B. (2026). Generative AI and firm values. *Journal of Finance*, Early View.

[21] Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2024). GPTs are GPTs: Labor market impact potential of LLMs. *Science*, 384(6702), 1306–1308.

[22] Engelberg, J., Manela, A., Mullins, W., & Vulicevic, L. (2025). Entity neutering (Working paper).

[23] Farber, H. S., & Gibbons, R. (1996). Learning and wage dynamics. *Quarterly Journal of Economics*, 111(4), 1007–1047.

[24] Felten, E., Raj, M., & Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, 42(12), 2195–2217.

[25] Gans, J. S., & Landry, P. (2026). Training your replacement: Worker incentives and the pace of automation (Working paper).

[26] Gao, J. (2026). Artificial intelligence, automation, and the allocation of effort. *American Finance Association Meetings*.

[27] Gao, Z., Jiang, W., & Yan, Y. (2025). A test of lookahead bias in LLM forecasts (Working paper).

[28] Gelbrich, M. (1990). On a formula for the L^2-Wasserstein metric between measures on Euclidean and Hilbert spaces. *Mathematische Nachrichten*, 147, 185–203.

[29] Glasserman, P., & Lin, C. (2023). Assessing look-ahead bias in stock return predictions generated by GPT sentiment analysis (Working paper).

[30] Grossman, S. J., & Stiglitz, J. E. (1980). On the impossibility of informationally efficient markets. *American Economic Review*, 70(3), 393–408.

[31] Hansen, A. L. (2026). Validating large language model annotations (feds 2026-020).

[32] Hartley, E. (2025). Narratives to numbers: Large language models and economic policy uncertainty (Working paper).

[33] Hassan, T. A., Hollander, S., van Lent, L., & Tahoun, A. (2019). Firm-level political risk: Measurement and effects. *Quarterly Journal of Economics*, 134(4), 2135–2202.

[34] Hassan, T. A., Kalyani, A., & Restrepo, P. (2026). The skill premium in times of rapid technological change (nber w34939).

[35] He, S., Lv, L., Manela, A., & Wu, J. (2025). Chronologically consistent large language models (Working paper).

[36] Humlum, A., & Vestergaard, E. (2025). Still waters, rapid currents: Early labor market transformation under generative AI (nber w33777).

[37] Iacoviello, M., & Tong, J. (2026). The AI-GPR index: Measuring geopolitical risk using artificial intelligence (Working paper).

[38] Jansen, M., Nagel, F., Yannelis, C., & Zhang, A. L. (2022). Data and welfare in credit markets (nber w30235).

[39] Jha, M., Qian, J., Weber, M., & Yang, B. (2024). ChatGPT and corporate policies (nber w32161).

[40] Jovanovic, B. (2014). Misallocation and growth. *American Economic Review*, 104(4), 1149–1171.

[41] Jovanovic, B., & Ma, S. (2023). Growth through learning. *Review of Economic Dynamics*, 50, 211–234.

[42] Jovanovic, B., & Rousseau, P. L. (2026). Artificial intelligence and task efficiency. *Journal of Monetary Economics*, 157.

[43] Kim, A. G., Muhn, M., & Nikolaev, V. (2023). From transcripts to insights: Uncovering corporate risks using generative AI (Working paper).

[44] Knott, M., & Smith, C. S. (1984). On the optimal mapping of distributions. *Journal of Optimization Theory and Applications*, 43(1), 39–49.

[45] Lang, M. H., & Lundholm, R. J. (1996). Corporate disclosure policy and analyst behavior. *The Accounting Review*, 71(4), 467–492.

[46] Lange, F. (2007). The speed of employer learning. *Journal of Labor Economics*, 25(1), 1–35.

[47] Levy, B. (2026). Caution ahead: Numerical reasoning and look-ahead bias in AI models. *Journal of Accounting Research*, 64(3), 1139–1188.

[48] Lindenlaub, I. (2017). Sorting multidimensional types: Theory and application. *Review of Economic Studies*, 84(2), 718–789.

[49] Lindenlaub, I., Oh, R., Rodriguez, M. A., & Veldkamp, L. (2026). Beyond exposure: Predicting AI adoption based on comparative advantage (nber w35271).

[50] Lise, J., & Postel-Vinay, F. (2020). Multidimensional skills, sorting, and human capital accumulation. *American Economic Review*, 110(8), 2328–2376.

[51] Lopez-Lira, A., Tang, Y., & Zhu, M. (2025). The memorization problem: Can we trust LLMs' economic forecasts? (Working paper).

[52] Ludwig, J., Mullainathan, S., & Rambachan, A. (2026). Large language models: An applied econometric framework. *Annual Review of Economics*, 18, 283–316.

[53] Sarkar, S. K., & Vafa, K. (2025). Lookahead bias in pretrained language models (Working paper).

[54] Sorensen, M. (2007). How smart is smart money? A two-sided matching model of venture capital. *Journal of Finance*, 62(6), 2725–2762.

[55] U.S. Census Bureau. (2026). Business Trends and Outlook Survey: Artificial intelligence supplement [Data set].

[56] Villani, C. (2003). *Topics in Optimal Transportation*. Providence, RI: American Mathematical Society.

[57] Wu, K., Yang, B., Ying, Z., & Zhou, D. (2025). Anonymization and information loss (Working paper).

[58] Zhang, X. F. (2006). Information uncertainty and stock returns. *Journal of Finance*, 61(1), 105–137.

Online Appendix

The Information Value of Artificial Intelligence: Evidence from Firm Growth and Labor Allocation

Boyan Jovanovic Sai Ma

A Data Construction

A.1 Capital IQ Transcripts and Classification Prompt

We construct the firm scores only from earnings calls dated before ChatGPT's release. For each call, we retain the Capital IQ company and transcript identifiers, the call date, and the text surrounding the relevant discussion, and we remove the firm name and ticker before classification. GPT-5.5 then assigns the five firm-level scores with the prompt reproduced below, which gives explicit definitions and a scoring scale but no examples of scored text. The temperature is set to zero for all earnings-call classifications, and document identifiers link each classification to its source transcript.

We also apply the same definitions to business descriptions for the first 400 firms in the management-guidance sample, ordered by Compustat identifier, after removing firm names and tickers. This separate GPT-5.5 classification uses 2021–2022 10-K filings. It does not use the earnings-call scores or any outcome data. After applying the confidence and outside-disclosure filters, 389 firms also appear in the earnings-call sample. The cross-document correlations and bootstrap intervals are reported in Section 3.3.1.

CLASSIFICATION PROMPT:

You are classifying pre-ChatGPT public-company earnings-call excerpts for an academic study of firms' internal information and production problems. Use only the supplied masked excerpt. Do not infer the company, use outside knowledge, or predict later success. Do not score whether an outside investor finds the disclosure readable. Score the operating decisions described by the firm.

Return integer scores from 0 to 100. Use 0–20 when the text supplies little direct evidence, 21–40 for a secondary activity, 41–60 for a material activity, 61–80 for a

central activity, and 81–100 only when repeated text shows that the activity dominates operations.

`internal_state_uncertainty`: How much the firm must estimate payoff-relevant states before committing resources. Examples include worker or project quality, customer credit, demand, diagnoses, claims, supplier reliability, inventory needs, and research outcomes. Generic business risk does not count.

`information_fragmentation`: How dispersed, qualitative, unstructured, or difficult to compare are the inputs used to estimate those states. Large data volumes alone do not count when a simple rule or structured database suffices.

`decision_stakes`: How economically important and difficult to reverse are the worker, project, capital, customer, supplier, or inventory decisions that rely on those estimates.

`genai_information_fit`: Holding product demand and production technology fixed, how much could language-model extraction, comparison, synthesis, or interpretation of existing text and records improve the precision of those estimates? Do not award points merely because software, data, or AI is named.

`direct_productivity_exposure`: Holding information precision fixed, how much could language generation, coding, summarization, or related generative tools reduce the labor or time required to produce the firm's goods or services?

`outside_disclosure_only`: Set to true when the excerpt supports difficulty for an outside reader but does not document a corresponding internal information or allocation problem. Otherwise set to false.

Give one concise evidence phrase for internal information exposure and one for direct productivity exposure. Confidence is 0 to 100 and measures whether the excerpt is sufficient for classification.

Classify every document in this JSON array:

```
[{"document_id": "...", "excerpt": "..."}, ...]
```

The classification records the following measures:

`document_id`	`internal_state_uncertainty`
`information_fragmentation`	`decision_stakes`
`genai_information_fit`	`direct_productivity_exposure`
`outside_disclosure_only`	`information_evidence`
`productivity_evidence`	`confidence`

Numeric measures are restricted to integers between zero and one hundred. We retain a document only when all required measures and its document identifier are present. We average eligible document scores within a firm-quarter and then average the quarter means across quarters.

Timing of GenAI disclosures. To date when firms first discuss their own GenAI activity in public, we use a separate set of 4,624 AI-related transcript excerpts from 2,483 Capital IQ companies between 2021 and 2025. A fixed search for generative AI, GenAI, ChatGPT, GPT, large language model, LLM, Copilot, Gemini, Claude, Llama, and foundation model identifies explicit generative-AI language and records the first occurrence for each company. Using fixed prompts, we retain passages in which firms describe their own GenAI products, operations, workforce, strategy, or investment with confidence of at least 0.70. These dates record public disclosure rather than operational adoption.

A.2 Other Data Sources

AI-related workers. The Department of Labor Office of Foreign Labor Certification provides Labor Condition Application records for H-1B, H-1B1, and E-3 workers. The files contain employer names, work locations, job titles, occupational codes, case status, intended employment dates, and certified wage rates. Firm names are harmonized and matched to Compustat parents. We remove legal suffixes, harmonize punctuation and abbreviations, and manually review large employers and ambiguous matches.

We classify an application as AI- or machine-learning-related from job titles and occupational codes. Within this set, information roles include data science, statistics, operations research, business intelligence, economics, and related decision-support occupations. Execution roles include software development, systems engineering, architecture, database administration, and related engineering occupations. The dictionaries are fixed before estimating firm outcomes. In Tables 3 and 6, prior AI hiring is an indicator for any AI- or machine-learning-related application filed before ChatGPT's release. Table A8 uses the log of one plus the number of these applications. The main worker regressions use 2018–2025 applications and annualized certified wages.

Capital IQ professional histories provide the broader hiring measure. We classify a job title as an information role if it contains any of the following terms: data, analytics, science, strategy, insight, economics, risk, planning, forecasting, research, intelligence, or quantita-

tive. Among the remaining titles, those containing engineering, technology, software, developer, digital, information officer, architect, infrastructure, or systems are execution roles. We use the first observed role of each person at a firm and collapse multiple titles that begin in the same year to one person-firm spell. When a person has both information and execution titles beginning in that first year, we classify the spell as an information spell. Titles containing none of these terms are excluded. The firm-year share of new hires in information roles is the number of information spells divided by the sum of information and execution spells. Using the first observed role excludes subsequent promotions and other internal role changes from the new-hire count.

Analyst-based opacity measure. Opacity is constructed from I/B/E/S analyst forecasts before 2022. For each firm and forecast horizon, we compute analyst disagreement, the range of forecasts, and realized forecast error after earnings are announced. Each component is standardized within this pre-2022 sample, and opacity is their average. A high value therefore means that analysts disagreed more, issued a wider range, or missed realized earnings by more in those years.

Analyst-based opacity differs from low analyst coverage. The number of analysts following a firm can reflect firm size, investor neglect, or limited institutional ownership. The opacity measure uses disagreement, range, and forecast errors because these quantities are closer to signal noise. Table 7 reports analyst coverage in a separate column.

Firm variables. Compustat provides annual sales, employment, assets, cost of goods sold, and operating expenses. We measure labor productivity as sales divided by employment. We winsorize annual growth rates before estimation. The specification for firm performance used in the main results removes the linear trend estimated from each firm's 2018–2022 observations and absorbs firm and two-digit industry-by-year fixed effects.

B Proofs, Specifications, and Model Details

B.1 Proof Details

Posterior precision and expected decision efficiency. Let $\theta \sim \mathcal{N}(\bar{\theta}, 1/p_0)$ and $s = \theta + \xi$, where $\xi \sim \mathcal{N}(0, 1/p_A)$ is independent of θ. Normal updating gives

$$v(p_A) = \frac{1}{p_0 + p_A}, \qquad \widehat{\theta}(s) = \bar{\theta} + q(p_A)(s - \bar{\theta}), \qquad q(p_A) = \frac{p_A}{p_0 + p_A}.$$

Maximizing expected decision efficiency is equivalent to minimizing expected squared decision error, so $d^* = \widehat{\theta}(s)$. If $Z = (\theta - \widehat{\theta}(s))/\sqrt{v(p_A)}$, then Z is standard normal and

$$e = 1 - [1 - q(p_A)]Z^2.$$

It follows that $E[e] = q(p_A)$ and $\text{var}(e) = 2[1 - q(p_A)]^2$. Moreover,

$$q'(p_A) = \frac{p_0}{(p_0 + p_A)^2} > 0, \qquad q''(p_A) = -\frac{2p_0}{(p_0 + p_A)^3} < 0,$$

so $q(0) = 0$ and $q'(0) = 1/p_0$.

Proof of Theorem 1. Write $x^W = X^W - \mu^W$ and $\widetilde{m}^F = m^F - \mu_m$. The terms in equation (2) that depend on assignment are $\widetilde{m}^{F\prime} x^W$. The equilibrium assignment therefore solves

$$\max_{\pi \in \mathcal{P}(\widetilde{m}^F, x^W)} E_\pi[\widetilde{m}^{F\prime} x^W],$$

where $\mathcal{P}(\widetilde{m}^F, x^W)$ contains the joint distributions with the given marginals. Because

$$E_\pi \|\widetilde{m}^F - x^W\|^2 = E\|\widetilde{m}^F\|^2 + E\|x^W\|^2 - 2E_\pi[\widetilde{m}^{F\prime} x^W],$$

this is the quadratic-cost transport problem studied in Villani (2003), here with Gaussian marginals. Knott and Smith (1984) and Gelbrich (1990) show that its unique solution is $x^W = B_t \widetilde{m}^F$, where B_t is given by equation (12). It preserves the worker distribution since

$$B_t \Sigma_{m,t} B_t = \Sigma_W.$$

The matrix B_t is symmetric and positive definite. Hence, for two distinct marginal-value vectors m_i^F and m_j^F,

$$(m_i^F - m_j^F)'[X^W(m_i^F) - X^W(m_j^F)] = (m_i^F - m_j^F)' B_t (m_i^F - m_j^F) > 0,$$

which establishes positive sorting in marginal values.

The wage schedule

$$W_t(X^W) = \bar{W}_t + \mu'_{m,t} x^W + \frac{1}{2}(x^W)' B_t^{-1} x^W$$

has gradient $\mu_{m,t} + B_t^{-1} x^W = m^F$ at the assigned worker and therefore supports the allocation. When q and g are common across firms, $m^F - \mu_m = A(q,g)' x^F$. Substitution gives $x^W = T_t(q,g) x^F$ with $T_t(q,g) = B_t A(q,g)'$. If $A(q,g)$, $\Sigma_{F,t}$, and Σ_W are diagonal with positive entries, $T_t(q,g)$ is diagonal with positive entries and sorting is positive within each task.

Proof of Proposition 1. At the prevailing wage schedule, the assignment map for a measure-zero firm is $X_i^W = \mu^W + B_t(m_i^F - \mu_{m,t})$. Since

$$\frac{dm_i^F}{dp_{A,i}} = \begin{pmatrix} 0 \\ a_I X_{I,i}^F q'(p_{A,i}) \end{pmatrix},$$

differentiating the assignment map gives equation (15). Differentiating equation (2) along this assignment gives equation (16). Finally, equation (4) implies that the change in next period's firm capabilities is $\Gamma\, dX_i^W/dp_{A,i}$. The stated sign conditions imply that output rises both with assignment fixed and through reassignment, and that next period's firm capabilities also rise.

Proof of Corollary 1. Holding assignment fixed, $\partial Q_i/\partial p_{A,i} = \ell_i q'(p_{A,i})$. Substituting the derivatives of q from the posterior calculation gives equation (17).

Proof of Proposition 2. With execution terms fixed and $c = 0$, the firm's first-order condition is $W'(X_I^W, q) = a_I q X_I^F$. The equilibrium assignment $X_I^W = \kappa X_I^F$ therefore implies $W'(X_I^W, q) = a_I q X_I^W/\kappa$. Integration gives equation (18), up to the outside payoff $\bar{W}$. Differentiating with respect to p_A gives the wage response in the proposition.

Proof of Corollary 2. Differentiate the wage response in Proposition 2 with respect to X_I^W.

Proof of Proposition 3. Differentiate equation (2) with respect to g_i, first holding the match fixed and then differentiating the worker's execution marginal product. Positive execution capabilities and skills give the signs in equation (19).

Proof of Lemma 1. Realized output differs from its expectation by $\epsilon_{q,i} = \ell_i[e - q(p_{A,i})] + u_i$, and $E[\epsilon_{q,i} \mid X_i^F, X_i^W] = 0$ because $E[e] = q(p_{A,i})$. Independence of u_i and e gives

$$\operatorname{var}(\epsilon_{q,i} \mid X_i^F, X_i^W) = 2\ell_i^2[1 - q(p_{A,i})]^2 + \sigma_u^2,$$

which is equation (20). If $\sigma_u^2 = 0$, then $e - q(p_{A,i}) = [1 - q(p_{A,i})](1 - Z^2)$ and

$$E|\epsilon_{q,i}| = |\ell_i|[1 - q(p_{A,i})]E|1 - Z^2| = 4\varphi(1)|\ell_i|[1 - q(p_{A,i})].$$

The last equality uses $E|1-Z^2| = 2\int_{-1}^{1}(1-z^2)\varphi(z)\,dz = 4\varphi(1)$. Moreover,

$$\frac{d}{dp_{A,i}}\operatorname{var}(\epsilon_{q,i} \mid X_i^F, X_i^W) = -4\ell_i^2[1-q(p_{A,i})]q'(p_{A,i}) < 0.$$

For a fixed $\Delta p_{A,i} > 0$, the decline in forecast-error variance is

$$2\ell_i^2\left\{[1-q(p_{A,i})]^2 - [1-q(p_{A,i}+\Delta p_{A,i})]^2\right\}.$$

The expression in braces is larger at lower initial precision, and the decline rises with $|\ell_i|$ through the factor $2\ell_i^2$.

Lemma 2 (Annual guidance errors) *Suppose firm i makes N_d decisions with precision $p_{A,i}$ during a reporting year. They receive equal weight, and their posterior errors are independent conditional on the firm. If annual output depends on average decision efficiency, then*

$$E\left|\ell_i[\overline{e} - q(p_{A,i})]\right| = \kappa_d(N_d)|\ell_i|[1-q(p_{A,i})],$$
$$\kappa_d(N_d) = 2\left[F_{\chi^2_{N_d}}(N_d) - F_{\chi^2_{N_d+2}}(N_d)\right], \tag{A1}$$

where $F_{\chi^2_k}$ is the distribution function of a chi-squared random variable with k degrees of freedom. In particular, $\kappa_d(1) = 4\varphi(1)$, and $\sqrt{N_d}\,\kappa_d(N_d)$ converges to $2/\sqrt{\pi}$ as N_d increases.

Proof. For decision j, write

$$e_j - q(p_{A,i}) = [1-q(p_{A,i})](1-Z_j^2), \qquad Z_j \sim \mathcal{N}(0,1).$$

Independence implies $C_{N_d} = \sum_{j=1}^{N_d} Z_j^2 \sim \chi^2_{N_d}$ and

$$\overline{e} - q(p_{A,i}) = [1-q(p_{A,i})]\left(1 - \frac{C_{N_d}}{N_d}\right).$$

Because $E[C_{N_d}] = N_d$,

$$E|C_{N_d} - N_d| = 2E[(N_d - C_{N_d})\mathbf{1}\{C_{N_d} \le N_d\}].$$

The chi-squared density gives $E[C_{N_d}\mathbf{1}\{C_{N_d} \le N_d\}] = N_d F_{\chi^2_{N_d+2}}(N_d)$. Dividing by N_d yields equation (A1). The expression for $\kappa_d(1)$ follows from Lemma 1, and the limiting expression follows from the normal approximation to a chi-squared random variable.

Proof of Proposition 4. Suppress firm and time subscripts and write a_q for the number of signals in the preceding period. Substituting equation (24) into equation (25), the part of the objective that depends on investment is

$$q\chi n^{\alpha} a_q^{1-\alpha} - n.$$

The first-order condition is

$$q\chi\alpha n^{\alpha-1} a_q^{1-\alpha} = 1.$$

Because $0 < \alpha < 1$, the objective is strictly concave in n and this condition gives the unique optimum. Dividing by a_q gives

$$\frac{n}{a_q} = (\alpha\chi q)^{1/(1-\alpha)},$$

which proves equation (26). Substitution into equation (24) gives equation (27). Both expressions increase in q. Finally, the other terms in equation (2) remain stationary once firm capabilities and assignment have converged. As a_q grows, qa_q therefore gives output per worker the limiting growth rate $G - 1$.

Opacity and the information channel. Hold the match (X_i^F, X_i^W) and production efficiency fixed. From equation (2),

$$\begin{aligned} &Q_i(X_i^F, X_i^W, p_A + \Delta p_A, g_i) - Q_i(X_i^F, X_i^W, p_A, g_i) \\ &\quad = [q(p_A + \Delta p_A) - q(p_A)] \left(a_{q,i} + a_I X_{I,i}^F X_{I,i}^W\right) \\ &\quad = \Delta Q_i(p_A, \Delta p_A). \end{aligned}$$

The assigned worker's skills before and after the increase in precision are $X_i^{W,0}$ and $X_i^{W,1}$. The total output change from the information channel decomposes as

$$\begin{aligned} &Q_i(X_i^F, X_i^{W,1}, p_{A,i} + \Delta p_{A,i}, g_i) - Q_i(X_i^F, X_i^{W,0}, p_{A,i}, g_i) \\ &= \underbrace{Q_i(X_i^F, X_i^{W,0}, p_{A,i} + \Delta p_{A,i}, g_i) - Q_i(X_i^F, X_i^{W,0}, p_{A,i}, g_i)}_{\text{assignment held fixed}} \\ &\quad + \underbrace{Q_i(X_i^F, X_i^{W,1}, p_{A,i} + \Delta p_{A,i}, g_i) - Q_i(X_i^F, X_i^{W,0}, p_{A,i} + \Delta p_{A,i}, g_i)}_{\text{reassignment}}. \end{aligned}$$

When $\ell_i = a_{q,i} + a_I X_{I,i}^F X_{I,i}^W > 0$, the increase in $q(p_{A,i})$ implies $\Delta Q_i > 0$. The estimated ℓ_i is positive by construction.

Expected decision efficiency is strictly concave in precision, so the increment $q(p_A + \Delta p_A) - q(p_A)$ is strictly decreasing in p_A: its derivative is $q'(p_A + \Delta p_A) - q'(p_A) < 0$ for $\Delta p_A > 0$.

For two firms with the same coefficient $\ell > 0$ and $p_{A,H} < p_{A,L}$, this gives $\Delta Q_H > \Delta Q_L$. The output gain with assignment held fixed also increases with $\Delta p_{A,i}$ and ℓ_i. These inequalities hold at the firm's initial match. The second term in the decomposition is the reassignment contribution.

Empirical specifications. The main regressions use the common specification

$$Y_{it} = \alpha_i + \delta_{s(i),t} + \beta_I Post_t \times I_i + \beta_P Post_t \times P_i + \varepsilon_{it},$$

where $s(i)$ denotes the relevant industry or guidance type for each dependent variable. Information channel exposure I_i and production channel exposure P_i enter every specification jointly.

For management guidance, the unit is a firm, fiscal target year, and quarter of guidance issuance. The guidance regression includes firm fixed effects and interacts guidance type with the quarter of issuance. Guidance type is defined by forecast horizon, point versus range guidance, and whether the guidance is an update. Standard errors are two-way clustered by firm and quarter of issuance. The benchmark forecasts use only firm and industry sales histories available before guidance is issued. Table A9 adds two-digit industry-by-issuance-quarter effects and varies the treatment of horizon and updates.

For Capital IQ, the unit is a firm-year. The regression of the share of new hires in information roles includes firm and year effects. The hire-count regressions use conditional Poisson with firm and two-digit industry-by-year effects, so their coefficients are proportional changes in expected hires and firm-years with zero hires remain in the sample. The H-1B application counts include firm and year effects. The wage regression is at the application level and adds state and occupation effects. Standard errors in the worker regressions are clustered by firm.

For firm performance, we estimate each firm's linear 2018–2022 trend in the dependent variable and subtract the projected trend from all available annual growth observations. The regression absorbs firm and two-digit industry-by-year fixed effects and clusters by firm. Sales, employment, and sales per worker use the same observations and transformation, so the reported coefficients satisfy the accounting identity among the three growth rates.

The supporting analyst-forecast regressions use firm-quarter observations, firm fixed effects, and two-digit industry-by-quarter fixed effects. The alternative specifications in Table A7 add pre-2023 macroeconomic exposures, omit 2021–2022, move the break to placebo

years, and separate information channel exposure from the other information characteristics. Table A8 replaces the two-digit industry effects with three- and four-digit industry effects and adds predetermined firm and technology characteristics.

B.2 The Production Channel

Section 2 derives the information channel from the precision of the signals a firm acts on. The production channel rests on a different primitive. The share of execution operations performed by generative AI tools provides a microfoundation for the production-efficiency increase g_i used in Section 5.1.

The execution task consists of a unit continuum of elementary operations. Without generative AI tools, each operation requires one unit of execution labor. The tools perform a share $\lambda_i \in [0, 1)$ of the firm's operations without labor, so the unit labor requirement falls to $1 - \lambda_i$ and output per unit of execution labor rises by the proportional amount

$$g_i = \frac{\lambda_i}{1 - \lambda_i}.$$

The production-efficiency increase g_i is zero when no operation is exposed, and it is increasing and convex in the exposed share because each operation the tools take over economizes on a labor requirement that is already smaller.

The production channel exposure measure ranks firms by the share of production work generative AI tools can carry out. The estimated specification maps production exposure into production efficiency according to

$$g_i \simeq \delta_g (P_i - P_{10})_+,$$

where P_i is expressed in standard-deviation units, as in Section 5.1, and δ_g is the production-efficiency gain per standard deviation of exposure above the tenth percentile. Placing the zero at the tenth percentile leaves the least exposed decile with no measured gain. Moving the floor tests how much the zero point matters for the decomposition. The information-channel share is 35.4 percent with the floor at the minimum and 54.4 percent with the floor at the twenty-fifth percentile in Table A20.

Output per unit of execution labor rises by the factor $1 + g_i$. The same factor multiplies the execution entry of $A(q, g_i)$ and therefore the firm's marginal value of worker execution skill. The information and production channels both change firms' marginal values of worker skills, and they interact in assignment whenever $c \neq 0$.

To map the production channel into observed sales and employment, equation (23) combines the direct gain g_i with the increase in the execution interaction at the initial match. Employment changes by $\eta_L \Delta Q_i^P$, and sales growth equals the sum of sales-per-worker growth and employment growth. Together, δ_g and η_L reproduce the empirical decomposition in which roughly one-third of the sales-per-worker growth coefficient reflects higher sales and two-thirds reflects lower employment. Because g_i is independent of signal precision, it does not generate the opacity interaction that distinguishes the information channel.

B.3 Cobb–Douglas Task Production

The diagonal interactions in equation (2) arise from Cobb–Douglas production within each task. Suppose the primitive firm and worker inputs in task $k \in \{I, E\}$ are $Z_{k,i}^F > 0$ and $Z_{k,i}^W > 0$. The information and execution contributions to match output are

$$q_i a_I (Z_{I,i}^F)^{\alpha_I} (Z_{I,i}^W)^{1-\alpha_I} + (1+g_i) a_E (Z_{E,i}^F)^{\alpha_E} (Z_{E,i}^W)^{1-\alpha_E},$$

where $0 < \alpha_I, \alpha_E < 1$. Express firm capabilities and worker skills as

$$X_{I,i}^F = (Z_{I,i}^F)^{\alpha_I}, \quad X_{I,i}^W = (Z_{I,i}^W)^{1-\alpha_I}, \quad X_{E,i}^F = (Z_{E,i}^F)^{\alpha_E}, \quad X_{E,i}^W = (Z_{E,i}^W)^{1-\alpha_E}.$$

The two terms then become $q_i a_I X_{I,i}^F X_{I,i}^W$ and $(1+g_i) a_E X_{E,i}^F X_{E,i}^W$, exactly the diagonal interactions in equation (2). The direct term $a_{q,i} q_i$ allows decision efficiency to affect output apart from worker-firm complementarity. Similarly, g_i records the direct production-efficiency gain, while $1 + g_i$ raises the value of execution capability and skill. The cross-task interaction c completes the more general specification used in the paper.

C Supporting Results

C.1 Measurement and Robustness

Validation across models, documents, and human coders. Table A1 asks whether the transcript measures depend on the language model, source document, or evaluator. Each panel changes one of these elements while retaining the same definitions.

Figure A1 puts the cross-document agreement in the original 0–100 scoring range. Firms in the top earnings-call quintile of information channel exposure score 18.4 points higher in their 10-K filings than firms in the bottom quintile. The difference is statistically significant with a p-value below 0.001. The corresponding difference for production channel exposure is

Table A1: Validation Across Models, Documents, and Human Coders

Sources	Measure	Correlation
Panel A. Alternative language models		
GPT-5.5 and Claude Opus 4.8	Information channel exposure	0.91
GPT-5.5 and Claude Sonnet 5	Information channel exposure	0.83
GPT-5.5 and Claude 3 Haiku	Information channel exposure	0.87
GPT-5.5 and Llama 3 8B	Information channel exposure	0.82
GPT-5.5 and Claude 3 Haiku	Production channel exposure	0.93
GPT-5.5 and Llama 3 8B	Production channel exposure	0.91
GPT-5.5 and Claude 3 Haiku	Adjusted information channel exposure	0.86
GPT-5.5 and Llama 3 8B	Adjusted information channel exposure	0.88
Panel B. Annual reports and earnings calls		
Annual reports and calls	Information channel exposure	0.40 [0.30, 0.49]
Annual reports and calls	Production channel exposure	0.86 [0.83, 0.89]
Annual reports and calls	Four information dimensions	0.50 [0.41, 0.58]
Annual reports and calls	Information channel exposure, conditional on production, size, and industry	0.33 [0.17, 0.46]
Panel C. Human and language-model scores		
Research assistant 1 and GPT-5.5	Information channel exposure	0.87
Research assistant 2 and GPT-5.5	Information channel exposure	0.89
Research assistant 1 and GPT-5.5	Production channel exposure	0.91
Research assistant 2 and GPT-5.5	Production channel exposure	0.79

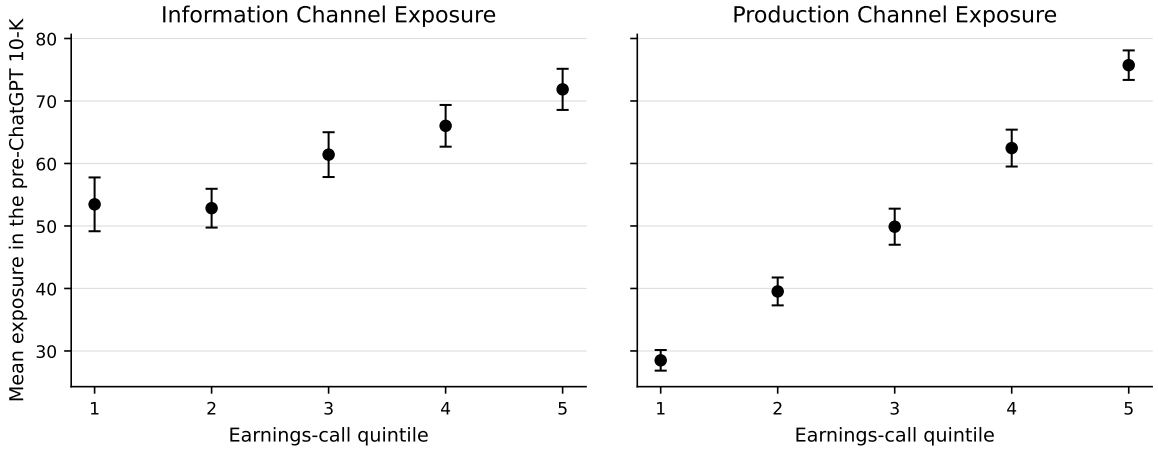
Notes. Panel A holds the classification prompt fixed. Panel B reports rank correlations for 389 firms scored from both 2021–2022 annual reports and earnings calls held before ChatGPT's release. The four information dimensions are state uncertainty, information fragmentation, decision stakes, and information channel exposure. The last row in Panel B accounts for production channel exposure, log assets, and detailed industry separately in each document before calculating the rank correlation. Brackets contain 95 percent bootstrap intervals. Panel C uses 100 randomly selected firms scored independently by two research assistants who did not observe the language-model scores.

47.2 points. Thus, even though production activities are described more similarly across the two documents, the earnings-call measure also preserves economically meaningful differences in information channel exposure.

Table A2 shows how information channel exposure differs within detailed industries. Each pair has similar production channel exposure but substantially different information channel exposure, and the 10-K measure gives the same ordering. Higher-exposure firms combine evidence that differs in form and meaning, whereas lower-exposure firms rely more on quantities, product specifications, or physical production. The pairs therefore illustrate the distinction between learning what to do and carrying out the chosen task.

Workforce exposure to generative AI. Eisfeldt, Schubert, Zhang, and Taska (2026) score each O*NET task statement with a language model, average the task scores within

Figure A1: Variation in Information and Production Channel Exposure Across Documents



Notes. Dots report mean exposure from independently scored 2021–2022 Form 10-K filings for each quintile of the earnings-call measure. Lines report 95 percent confidence intervals. The sample contains 389 firms with eligible scores from both sources.

Table A2: Information Channel Exposure Within Industries

Industry	Higher information channel exposure	Lower information channel exposure	Call I	$\lvert\Delta P\rvert$	10-K I
Pharmaceuticals	BioMarin. Clinical, regulatory, safety, efficacy, and market evidence	Neogen. Safety tests, reagents, vaccines, and instruments	90 / 15	1.3	84 / 55
Clothing stores	Burlington. Changing vendor offers, fashion trends, and local demand	Abercrombie & Fitch. Own-brand forecasts, product age, and markdowns	71 / 7	0.8	67 / 43
Aerospace products	General Dynamics. Mission needs, technical requirements, workforce, and suppliers	TransDigm. Specified aircraft components and aftermarket demand	96 / 20	4.5	73 / 66

Notes. Rows pair firms in the same detailed industry. Call I and 10-K I list the two firms' information channel exposure, with the higher-exposure firm first. $\lvert\Delta P\rvert$ is the absolute difference in call-based production channel exposure. Scores range from 0 to 100. Reported pairs differ by at least three standard deviations in call-based information channel exposure, by no more than 0.4 standard deviation in call-based production channel exposure, and have the same 10-K ordering. The table reports three illustrations from 18 qualifying pairs across six industries. Business descriptions are based on the 10-K filings.

occupations, and weight occupations by each firm's employment shares in March 2022. Their measure covers 2,684 of the 4,477 scored firms. Panel A of Table A3 relates it to the two transcript measures, and Panel B re-estimates three baseline specifications on the firms it covers, first without and then with its interaction with Post.

Table A3: Transcript Measures and Workforce Exposure to Generative AI

Panel A. Workforce exposure and the transcript measures

	No industry effects (1)	Two-digit effects (2)	Four-digit effects (3)
Information channel exposure, I	-0.055^{***}	-0.016	-0.007
	(0.017)	(0.016)	(0.016)
Production channel exposure, P	0.330^{***}	0.138^{***}	0.064^{***}
	(0.018)	(0.021)	(0.023)
p-value: $I = P$	< 0.001	< 0.001	0.023
Observations	2,684	2,684	2,192

Panel B. Baseline specifications with workforce exposure

	Guidance error		Share of new hires in information roles		Sales-per-worker growth	
	Matched sample (1)	Add workforce exposure (2)	Matched sample (3)	Add workforce exposure (4)	Matched sample (5)	Add workforce exposure (6)
Post $\times I$	-1.344^{***}	-1.301^{***}	0.118^{***}	0.120^{***}	2.145^{***}	2.138^{***}
	(0.273)	(0.275)	(0.015)	(0.015)	(0.491)	(0.492)
Post $\times P$	-0.101	-0.239	-0.040^{**}	-0.056^{**}	1.894^{***}	1.954^{***}
	(0.270)	(0.366)	(0.020)	(0.022)	(0.708)	(0.716)
Post $\times$ workforce exposure	–	0.379	–	0.036	–	-0.390
	–	(0.614)	–	(0.027)	–	(0.719)
Observations	4,140	4,140	1,273	1,273	17,203	17,203
Firms	662	662	452	452	2,352	2,352

Notes. Workforce exposure is the firm-level GenAI exposure measure of Eisfeldt, Schubert, Zhang, and Taska (2026). Panel A regresses workforce exposure, standardized across scored firms, on I and P jointly. Columns (2) and (3) include two- and four-digit industry effects. Panel B uses observations with workforce-exposure data. For each dependent variable, the first column reports the baseline specification on this sample and the second adds Post interacted with workforce exposure. Coefficients are for a one-standard-deviation increase in each exposure measure. Fixed effects and standard errors follow column (1) of Table 4, column (1) of Panel A in Table 5, and column (1) of Panel B in Table 6. Panel A reports heteroskedasticity-robust standard errors. $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

Production channel exposure and Census AI use. Production channel exposure and workforce exposure both describe tasks that AI could affect. The Business Trends and Outlook Survey instead records the share of businesses reporting AI use. Table A4 relates that share to sector means of the two exposure measures across all available survey waves. Because the question wording changes during the sample, the last two rows report the estimates separately under each wording.

Table A4: Production Channel Exposure and Census AI Use

	P_i	Workforce exposure	p-value: equal	Waves	Observations
Current use, all waves	4.861***	0.611	0.013	69	1,076
	(0.908)	(0.758)			
Expected use, next six months	5.857***	0.769	0.013	69	1,097
	(1.036)	(0.896)			
Current use, production wording	3.893***	0.509	0.042	54	819
	(0.921)	(0.735)			
Current use, business-function wording	7.734***	1.346	0.009	15	257
	(0.999)	(1.129)			

Notes. The dependent variable is the percentage of businesses in a Census sector reporting current or expected AI use in the Business Trends and Outlook Survey. The sample contains 18 sectors and covers 69 survey waves from 2023 through 2026. The first 54 waves ask about AI use in producing goods or services, while the final 15 ask about use in any business function. The production channel exposure and workforce exposure measures are sector means formed on their common firm sample and standardized across sectors. Both measures enter jointly with survey-period effects. Coefficients are percentage points per sector standard deviation. Parentheses contain standard errors clustered by Census sector. The reported p-values and significance stars use exact restricted wild-cluster bootstrap inference across the 18 sectors. $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

Pretrend tests. The joint tests use the 2017–2021 coefficients for guidance and the 2018–2021 coefficients for hiring and sales-per-worker growth in Figure 1. All three tests cluster by firm and omit 2022. Panel C uses unadjusted annual growth, so its test is computed before subtracting firm trends. Figure 1 reports the p-values. Table A7 complements these tests with placebo dates and estimates excluding 2021–2022.

GenAI use in financial reporting. We use an external authorship measure to ask whether firms with greater information channel exposure subsequently use generative AI in financial reporting. Panel A of Table A5 uses changes after ChatGPT's release in the GenAI-authorship probabilities of Blankespoor, deHaan, and Li (2026). The baseline specification in column (2) removes text carried over unchanged from an earlier report. Columns (3) and (4) use adjusted information channel exposure and exclude conference calls, respectively. Panel B reports estimates for two earlier placebo periods and asks whether the same firms use generative AI across a broader set of financial reports. The information channel exposure coefficient is essentially zero in both placebo windows and positive afterward, whereas the production channel exposure coefficient is positive only in the second placebo window. The positive relation after ChatGPT's release is therefore specific to information channel exposure.

Table A5: Information Channel Exposure and GenAI Use in Financial Reports

Panel A. Use after ChatGPT's release	Complete report text (1)	New report text (2)	Adjusted Information Channel Exposure (3)	Exclude conference calls (4)
I	0.007*	0.010**	0.012***	0.012**
	(0.004)	(0.004)	(0.004)	(0.005)
P	−0.007	−0.008	−0.008	−0.005
	(0.006)	(0.005)	(0.005)	(0.006)
p-value: $I = P$	0.073	0.018	0.010	0.067
Observations	7,158	7,127	7,127	5,694
Firms	2,616	2,579	2,579	2,543
Panel B. Timing and use across reports	2018–2020 placebo (1)	2020–Nov. 2022 placebo (2)	Share of report types with high use (3)	At least two report types with high use (4)
I	−0.000	−0.000	0.013***	0.020***
	(0.005)	(0.005)	(0.005)	(0.007)
P	0.002	0.013**	−0.011*	−0.008
	(0.006)	(0.006)	(0.006)	(0.009)
p-value: $I = P$	0.848	0.117	0.008	0.029
Observations	7,292	7,101	2,579	2,338
Firms	2,420	2,476	2,579	2,338

Notes. The dependent variable follows the extensive-margin construction in Blankespoor, deHaan, and Li (2026). A firm-report pair is classified as high use if its change in the authors' GenAI-authorship score ever enters the top tercile for that report type and as low use if it always remains in the bottom tercile. Middle-tercile cases are omitted. It is an indicator, so a coefficient of 0.010 corresponds to one percentage point. New report text excludes unchanged passages. Panel A, column (3), replaces I with adjusted information channel exposure, and column (4) excludes conference calls. Columns (1)–(4) include I and P jointly, state uncertainty, log assets, two-digit industry effects, and report-type effects when applicable. Panel B, columns (3) and (4), aggregate the high-use classification across report types at the firm level. Standard errors are clustered by firm except in these two firm-level specifications, which use heteroskedasticity-robust standard errors. $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

GenAI use and guidance accuracy. Information channel exposure is constructed from earnings calls held before ChatGPT's release and does not record whether firms later use generative AI. Table A6 therefore asks whether the decline in management guidance error from 2023 onward is larger among firms whose financial reports show high GenAI use. Each column interacts Post $\times I$ and Post $\times P$ with the share of a firm's classified report types with high use, the firm-level measure in Panel B, column (3), of Table A5. In column (1), the decline in absolute midpoint error becomes larger as this share rises. The interaction remains negative and statistically significant when column (2) adds interactions of Post with log assets and the number of classified report types, and it is significant at the 10 percent level when column (3) excludes firms without a high-use report type. Column (4) places a placebo break in 2021 among guidance issued before 2023, and the share has no comparable relation with guidance error in that period. Thus, the decline in guidance error grows with

observed use in financial reporting, and this pattern is absent under the 2021 placebo break.

Table A6: GenAI Use in Financial Reports and Management Guidance Accuracy

	Baseline (1)	Size and coverage controls (2)	Firms with a high-use report type (3)	Placebo break in 2021 (4)
Post $\times I$	0.310	−0.269	0.114	0.444
	(0.585)	(0.651)	(0.885)	(0.659)
Post $\times I\times$ high-use share	−2.494***	−2.012**	−2.268*	−0.296
	(0.862)	(0.888)	(1.202)	(0.908)
Post $\times P$	0.121	0.353	−0.103	0.635
	(0.522)	(0.545)	(0.828)	(0.798)
Post $\times P\times$ high-use share	−0.170	−0.397	0.079	−0.670
	(0.721)	(0.678)	(1.118)	(0.983)
Post $\times I$ at a high-use share of one	−2.184***	−2.281***	−2.154***	0.148
	(0.398)	(0.398)	(0.428)	(0.355)
Observations	4,890	4,890	4,737	2,881
Firms	839	839	810	654

Notes. The dependent variable is absolute midpoint error in annual revenue guidance, in percentage points. High-use share is the share of a firm's classified report types with high use, the measure in Panel B, column (3), of Table A5. The last coefficient row adds the coefficients on Post $\times I$ and Post $\times I\times$ high-use share. Column (2) interacts Post, Post $\times I$, and Post $\times P$ with log assets and with the number of classified report types, expressed as deviations from their sample means. Column (3) excludes firms with no high-use report type. Column (4) uses guidance issued before 2023 and sets Post to one from 2021. All columns include Post × high-use share, firm FE, and guidance type × quarter FE. Standard errors are two-way clustered by firm and quarter. $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

Alternative specifications. Table A7 asks whether macroeconomic exposures, the 2021–2022 period, or other characteristics of the firm's information problem explain the post-2022 coefficients. Each column changes one element of the specification for management guidance error in Panel A and for sales-per-worker growth in Panel B.

Industry composition and predetermined characteristics. Firms with different exposure may also have faced different post-2022 industry conditions. Table A8 progressively narrows the industry variation and adds predetermined firm and technology characteristics. Its final panel asks whether permuting exposure across firms within four-digit industries reproduces the guidance coefficient.

Guidance timing and issuance. Table A9 examines the timing and composition of annual management guidance. Panel A absorbs industry-specific quarterly changes in forecastability and then varies the treatment of forecast horizon and within-year updates. Panel

Table A7: Alternative Specifications for Guidance Accuracy and Firm Performance

Panel A. Management revenue-guidance error

	Baseline (1)	Macro controls (2)	Exclude 2021–22 (3)	Placebo 2019 (4)	Add *U-F-K* composite (5)
Post $\times I$	−1.617***	−1.676***	−1.396***	0.207	−1.617***
	(0.271)	(0.261)	(0.338)	(0.294)	(0.271)
Post $\times P$	−0.051	−0.061	−0.122	0.233	−0.051
	(0.282)	(0.284)	(0.254)	(0.258)	(0.282)
Post × *U-F-K* composite	–	–	–	–	−0.007
	–	–	–	–	(0.156)
Observations	5,041	5,041	3,822	1,808	5,041
Firms	869	869	828	552	869
Firm and guidance type × quarter FE	Yes	Yes	Yes	Yes	Yes

Panel B. Sales-per-worker growth

	Baseline (1)	Macro controls (2)	Exclude 2021–22 (3)	Placebo 2016 (4)	Add *U-F-K* composite (5)
Post $\times I$	2.074***	1.522***	1.873***	−0.131	2.072***
	(0.501)	(0.383)	(0.246)	(0.239)	(0.500)
Post $\times P$	1.528**	1.932***	1.708***	−0.020	1.538**
	(0.721)	(0.530)	(0.340)	(0.275)	(0.725)
Post × *U-F-K* composite	–	–	–	–	−0.097
	–	–	–	–	(0.591)
Observations	22,236	22,236	17,643	21,792	22,236
Firms	3,109	3,109	3,559	3,253	3,109
Firm and two-digit industry × year FE	Yes	Yes	Yes	Yes	Yes

Panel C. Management guidance around the 2022 tightening cycle

	April–November 2022	2023 onward
Information channel exposure	0.359	−1.596***
	(0.440)	(0.280)
Production channel exposure	−0.609	−0.088
	(0.639)	(0.289)
Observations	5,041	
Firms	869	

Notes. Panel A uses management revenue-guidance error and Panel B sales-per-worker growth. Macro controls interact Post with pandemic sales volatility and opacity. Panel A also controls for 2021–22 guidance error, while Panel B controls for 2020–22 over-hiring and 2019–22 sales-per-worker growth. Column (3) drops 2021–22 and, in Panel B, uses unadjusted growth. Column (4) uses placebo dates of 2019 and 2016. Column (5) adds the *U-F-K* composite, the geometric mean of uncertainty, fragmentation, and stakes expressed in standard-deviation units. Panel C includes separate indicators for April 1 through November 29, 2022, and for 2023 onward in the baseline guidance specification. Standard errors are clustered by firm and quarter in Panels A and C and by firm in Panel B. $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

Table A8: Detailed Industry Effects and Predetermined Controls

			Common sample			
	Two-digit (1)	Three-digit (2)	Two-digit (3)	Four-digit (4)	Four-digit + firm characteristics (5)	Technology characteristics (6)
Panel A. Management revenue-guidance error						
Post $\times I$	-2.001^{***}	-2.594^{***}	-2.266^{***}	-2.702^{***}	-2.763^{***}	-2.037^{***}
	(0.425)	(0.464)	(0.456)	(0.531)	(0.536)	(0.429)
Post $\times P$	0.375	0.441	0.602	0.568	0.653	0.584
	(0.551)	(0.701)	(0.524)	(1.032)	(1.056)	(0.755)
Observations	3,595	3,250	2,744	2,744	2,744	3,595
Firms	629	577	502	502	502	629
Panel B. Sales-per-worker growth						
Post $\times I$	2.074^{***}	1.508^{***}	2.079^{***}	0.949^{*}	1.305^{***}	1.841^{***}
	(0.501)	(0.497)	(0.508)	(0.513)	(0.362)	(0.493)
Post $\times P$	1.528^{**}	1.975^{**}	1.481^{**}	0.895	1.648^{***}	1.468^{*}
	(0.721)	(0.792)	(0.732)	(0.815)	(0.553)	(0.789)
Observations	22,236	22,128	21,769	21,769	21,769	22,236
Firms	3,109	3,101	3,063	3,063	3,063	3,109

Panel C. Within-industry permutation of exposure	Actual coefficient	Mean permuted coefficient	At least as negative
Post $\times I$	-2.702	-0.005	0 of 400

Notes. Panel A assigns each firm its modal Compustat NAICS code across years, which is available for 3,595 guidance observations. The guidance-timing specifications use prior-year Compustat NAICS codes and retain 4,888 observations. All columns retain firm and guidance type × issuance-quarter effects and add the indicated NAICS industry × issuance-quarter effects. Panel B includes firm and the indicated NAICS industry × year effects. Columns (3)–(5) use the observations available for the four-digit specification. Column (5) interacts Post with pre-2023 leverage, cost structure, size, R&D, market-to-book, pandemic-era sales changes, over-hiring, and prior sales-per-worker growth. Column (6) interacts Post with industry AI exposure, prior AI hiring, and R&D intensity. Standard errors are clustered by firm and issuance quarter in Panel A and by firm in Panel B. $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$. Panel C uses the 2,744 observations and 502 firms in the four-digit specification in Panel A, with each firm's modal four-digit Compustat NAICS code. Each of 400 permutations moves I and P jointly across firms within four-digit industries and re-estimates the specification in column (4). The last column counts permutations with an information channel exposure coefficient at least as negative as the coefficient in the data. The finite-draw p-value is below 0.0025.

B asks whether firms with greater information channel exposure change the frequency with which they issue guidance.

Guidance error and mean reversion. Firms with larger guidance errors before 2023 have more room to improve. Table A10 therefore collapses the data to firm means and conditions each firm's mean error from 2023 onward on its mean error in earlier years. The remaining rows of Panel A vary how flexibly the specification absorbs the distribution of earlier errors, and Panel B repeats the specification for two earlier periods.

Table A9: Management Guidance: Timing and Issuance

Panel A. Absolute management-guidance error	Industry × quarter FE (1)	Add continuous horizon (2)	Initial annual guidance (3)
Post $\times I$	-1.621^{***}	-1.626^{***}	-1.612^{***}
	(0.317)	(0.318)	(0.318)
Post $\times P$	0.387	0.395	0.349
	(0.503)	(0.503)	(0.511)
Observations	4,888	4,888	4,660
Firms	855	855	846

Panel B. Guidance issuance	Any annual guidance (1)	log(1 + guidance count) (2)
Post $\times I$	0.003	0.003
	(0.004)	(0.003)
Post $\times P$	0.002	0.001
	(0.005)	(0.004)
Observations	35,816	35,816
Firms	4,477	4,477

Notes. Panel A retains the firm and guidance type × quarter effects from Table 4 and adds two-digit industry-by-issuance-quarter effects. Column (2) also controls for days to the fiscal year end. Column (3) retains the first annual guidance observation for each firm, target year, and accounting type. Panel B uses an annual panel of scored firms with firm and two-digit industry-by-year effects. Standard errors are clustered as in the corresponding main specification. $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

Adjusted information channel exposure. Information channel exposure I_i isolates language-model fit, whereas the adjusted measure also incorporates uncertainty, fragmentation, and stakes. Table A11 shows that replacing I_i with the adjusted measure preserves the signs across guidance, analyst forecasts, labor outcomes, and sales-per-worker growth.

Analyst-based opacity. Information channel exposure could instead proxy for the firm's initial information environment. Table A12 first enters analyst disagreement separately and then asks whether broader opacity amplifies the post-2022 changes in guidance, hiring, and sales-per-worker growth associated with information channel exposure.

Industry composition and worker measurement. Table A13 examines industry composition and alternative treatments of the worker data. Panel A excludes information and professional services, while Panel B varies the treatment of sparse Capital IQ hiring records and the concentration of H-1B applications among large filers. Panel C reports wage responses separately for information and execution hires. These coefficients motivate the sign

Table A10: Guidance Accuracy and Pre-2023 Error Levels

Specification	Information Channel Exposure	Firms
Panel A. Guidance error from 2023 onward		
Change in mean error	-2.160^{***}	699
	(0.357)	
Control for pre-2023 mean error	-1.618^{***}	699
	(0.249)	
Cubic controls for pre-2023 error	-1.628^{***}	699
	(0.253)	
Pre-2023 error-decile effects	-2.006^{***}	699
	(0.298)	
Trim pre-2023 error tails	-1.595^{***}	629
	(0.274)	
At least three pre-2023 years	-1.320^{***}	495
	(0.246)	
Panel B. Placebo before ChatGPT's release		
2020–2022 versus 2017–2019	0.057	537
	(0.266)	

Notes. The unit of observation is the firm. Annual revenue-guidance errors are averaged before 2023 and from 2023 onward. The change and decile specifications use the change in mean error. The other specifications in Panel A regress the later mean error on the indicated function of the earlier mean error. The trimmed sample removes the top and bottom five percent of earlier errors. Panel B applies the mean-error specification to two earlier periods. All specifications include production channel exposure, log assets, and two-digit industry effects. Errors are in percentage points. Heteroskedasticity-robust standard errors are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

Table A11: Responses to Adjusted Information Channel Exposure after 2022

	Management-guidance error (1)	Opacity (2)	Share of new hires in information roles (3)	Information-role applications (4)	Information-hire wage (5)	Sales-per-worker growth (6)
Post $\times I^{adj}$	-1.471^{***}	-0.039^{***}	0.098^{***}	0.112^{***}	0.014^{***}	1.660^{***}
	(0.228)	(0.010)	(0.014)	(0.020)	(0.005)	(0.533)
Post $\times P$	-0.199	0.019	-0.015	-0.112^{**}	-0.009	1.750^{**}
	(0.277)	(0.012)	(0.017)	(0.044)	(0.009)	(0.724)
Observations	5,041	97,513	1,756	1,824	25,970	22,236
Firms	869	4,051	635	228	223	3,109

Notes. Adjusted information channel exposure, I^{adj}, is the geometric mean of the uncertainty, fragmentation, stakes, and information channel exposure scores, standardized across firms. Each column follows the corresponding main-table specification. Columns (3)–(5) use Capital IQ new hires and H-1B applications and wages. Opacity averages standardized $\log(1 + x)$ disagreement, range, and absolute error. $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

Table A12: Information Channel Exposure and Analyst-Based Opacity

Panel A. Analyst disagreement	Management-guidance error (1)	Share of new hires in information roles (2)	Sales-per-worker growth (3)
Post $\times I$	-1.595^{***}	0.104^{***}	2.119^{***}
	(0.271)	(0.013)	(0.510)
Post $\times$ analyst disagreement	0.118	0.005	-0.624
	(0.352)	(0.027)	(0.933)
Observations	4,992	1,715	21,231
Firms	858	621	2,965
Panel B. Analyst-based opacity			
Post $\times I$	-2.095^{***}	0.113^{***}	2.251^{***}
	(0.319)	(0.015)	(0.533)
Post $\times$ opacity	-0.182	0.027	-0.860
	(0.264)	(0.021)	(0.762)
Post $\times I\times$ opacity	-0.953^{***}	0.015	1.834^{***}
	(0.195)	(0.014)	(0.557)
Post $\times P$	-0.007	-0.023	1.572^{**}
	(0.284)	(0.017)	(0.716)
Observations	5,041	1,756	22,236
Firms	869	635	3,109
Firm FE	Yes	Yes	Yes
Outcome-specific time FE	Yes	Yes	Yes

Notes. Panel A adds standardized analyst disagreement, measured before 2022, to the baseline specification. The P coefficient is included but not shown. Panel B estimates the displayed terms jointly. Opacity averages standardized disagreement, forecast range, and realized error over the same period. Columns follow the fixed effects, trend treatment, and clustering of the corresponding main-table specifications. Post begins in 2023. $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

restriction on c in Section 5.3.

The hire counts cover the same 1,215 firms in every year. The median firm-year has no classified hire, and the annual share of firms with at least one classified hire ranges from 14 to 32 percent. Panel B of Table A13 shows that the shift toward information roles remains when observations are weighted by the number of hires, when zero counts enter directly, and when the most recent years are excluded.

Margins and costs. Margins and operating income provide additional evidence on the production channel. In Table A14, production channel exposure has the larger relation with gross margins and is the only exposure measure that significantly predicts operating income per worker.

Table A13: Industry Composition and Labor Measurement

Panel A. Excluding information and professional services

	Management guidance error (1)	Sales-per-worker growth (2)
Post $\times I$	-1.850^{***}	2.319^{***}
	(0.337)	(0.594)
Post $\times P$	-0.110	1.504^{*}
	(0.432)	(0.784)
Observations	3,901	19,176
Firms	685	2,666

Panel B. Additional labor specifications

	Share of new hires in information roles, weighted (1)	Log information-role hires (2)	Log execution-role hires (3)	Share of new hires in information roles through 2024 (4)	Share of new hires in information roles through 2023 (5)	H-1B wage, full sample (6)	H-1B wage, drop top filers (7)
Post $\times I$	0.109^{***}	0.017^{***}	-0.019^{***}	0.089^{***}	0.093^{***}	0.012^{**}	0.014^{***}
	(0.013)	(0.004)	(0.005)	(0.016)	(0.021)	(0.005)	(0.005)
Post $\times P$	-0.049^{**}	-0.008^{*}	0.014^{*}	-0.046^{*}	-0.085^{*}	-0.005	-0.014
	(0.021)	(0.005)	(0.007)	(0.026)	(0.044)	(0.010)	(0.010)
Observations	1,756	9,720	9,720	1,419	1,156	25,923	13,067
Firms	635	1,215	1,215	534	454	223	213

Panel C. H-1B wages by task of hire

	Information hires (1)	Execution hires (2)
Post $\times I$	0.014^{***}	0.014^{*}
	(0.005)	(0.007)
Post $\times P$	-0.007	-0.013
	(0.009)	(0.015)
p-value: $I = P$	0.045	0.153
Observations	25,970	18,114
Firms	223	160

Notes. Panel A excludes NAICS 51 and 54 and otherwise follows the corresponding main specifications. Panel B, column (1), uses the same firm-years as Table 5, Panel A, column (1), weights each observation by its number of classified hires, and replaces year effects with two-digit industry-by-year effects. Columns (2) and (3) include zero-hire years for firms with a classified pre-2023 hire. Columns (4) and (5) end the sample in 2024 and 2023. The ten largest filers are excluded from column (7). Columns (2)–(5) include firm and two-digit industry-by-year effects. Columns (6) and (7) include firm, state, and occupation-by-year effects. Both columns of Panel C include firm, year, state, and occupation effects. Standard errors are clustered by firm. $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

Table A14: Margins, Costs, and Industry Composition

Panel A. Margins and costs

	Gross margin (1)	Operating expense/sales (2)
Post $\times I$	1.264***	−2.538***
	(0.364)	(0.522)
Post $\times P$	2.812***	−3.597***
	(0.451)	(0.678)
Pre-2023 mean	34.857	97.990
Observations	19,615	19,615
Firms	2,954	2,954

Panel B. Operating income and narrower industry exclusions

	Operating income per worker (1)	Sales per worker, narrower exclusions (2)
Post $\times I$	2.546	2.197***
	(2.053)	(0.665)
Post $\times P$	5.573**	1.821**
	(2.363)	(0.820)
Observations	10,857	18,151
Firms	1,675	2,526

Notes. Panel A reports ratios in percentage points. Gross margin is sales less cost of goods sold, divided by sales. The retained ranges are $[-100, 100]$ for gross margin and $[0, 300]$ for operating expense relative to sales. Panel B, column (1), reports pretrend-adjusted growth in operating income per worker. Column (2) excludes four-digit NAICS industries 5182, 5191, 5412, and 5615. All regressions include firm and two-digit industry-by-year effects. Standard errors are clustered by firm. $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

Upper-tail sales-per-worker growth. The long-run extension predicts that the information channel can expand the upper tail of firm growth. Table A15 reports both the probability of crossing a threshold in the pre-2023 growth distribution and the amount by which growth exceeds it.

C.2 Estimation and Numerical Implementation

Estimation details. Because the empirical moments come from different datasets, the numerical implementation preserves each moment's estimation sample and the covariance among estimates from overlapping firms. Management forecast moments come from LSEG guidance, assignment moments combine the pre-2023 relation between the share of Capital IQ new hires in information roles and information fragmentation with the H-1B wage response for information hires, and production-channel moments come from Compustat. The underlying information-fragmentation and production channel exposure scores are rescaled to lie between zero and one and measure $X^F_{I,i}$ and $X^F_{E,i}$, as described in Section 5.1. Decision stakes enter ℓ_i, the marginal contribution of decision efficiency to output, rather than the marginal values m^F_i that determine assignment. The assigned worker's skills before and after the increases in precision and production efficiency are $X^{W,0}_i$ and $X^{W,1}_i$, with centered values

Table A15: Information Channel Exposure and the Upper Tail of Sales-per-Worker Growth

	Threshold in the pre-2023 distribution		
	75th percentile	90th percentile	95th percentile
Panel A. Probability of exceeding the threshold			
Post $\times I$	0.0238***	0.0125***	0.0054*
	(0.0071)	(0.0043)	(0.0028)
Post $\times P$	0.0049	0.0034	0.0021
	(0.0098)	(0.0063)	(0.0041)
Panel B. Growth above the threshold			
Post $\times I$	0.5369***	0.2673**	0.1400*
	(0.1831)	(0.1171)	(0.0810)
Post $\times P$	0.2683	0.2601	0.2278*
	(0.2812)	(0.1882)	(0.1372)
Observations	22,236	22,236	22,236
Firms	3,109	3,109	3,109

Notes. Thresholds are fixed at the indicated quantiles of the pre-2023 distribution of detrended sales-per-worker growth. Panel A uses an indicator for growth above each threshold. Panel B uses the amount by which growth exceeds the threshold, with zero below it. All specifications include firm and two-digit industry-by-year fixed effects and remove each firm's 2018–2022 linear trend. Standard errors are clustered by firm. $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

$x_i^{W,s} = X_i^{W,s} - \mu^W$ for $s \in \{0, 1\}$. The numerical counterfactual measures the proportional change in output per worker as

$$\Delta Q_i = \ell_i \left[q(p_{A,i} + \Delta p_{A,i}) - q(p_{A,i})\right] + g_i + \bar{\ell} g_i X_{E,i}^F X_{E,i}^{W,0} + \bar{\ell}(m_i^{F,1} - \overline{m}^{F,1})'(x_i^{W,1} - x_i^{W,0}).$$

Here $\overline{m}^{F,1}$ is the employment-weighted mean of the marginal values $m_i^{F,1}$, and $X_{E,i}^{W,0} = \mu_E^W + x_{E,i}^{W,0}$ is the execution skill of the worker at the initial match, in levels. The coefficient ℓ_i measures $\partial Q_i / \partial q_i$ at the initial match, using information fragmentation and decision stakes. The third term is the change in the execution interaction implied by equation (2). Setting $g_i = 0$ isolates the information channel. Within that channel, holding assignment fixed sets $x_i^{W,1} = x_i^{W,0}$. The assignment map in Theorem 1 determines the reassignment term. Firm types receive mass proportional to their mean employment from 2019 through 2022 in both the assignment distribution and the output calculation. Centering $m_i^{F,1}$ with these same masses removes the component of worker productivity valued equally across firms, which does not create aggregate match surplus when workers are reassigned. The common multiplier $\bar{\ell}$ expresses both interaction terms as proportional changes in output.

Let L_i denote employment. It responds to the full production-channel gain in equation (23):

$$\Delta \log L_i = \eta_L \Delta Q_i^P.$$

Because sales-per-worker growth equals sales growth minus employment growth, the implied sales response is $(1+\eta_L)\Delta Q_i^P$. This accounting relation separates the change in workforce skill composition determined by assignment from the change in total employment.

To map the assigned worker's latent information skill $x_{I,i}^W$ into the share of new hires in information roles, we use the threshold rule $h_i = \Phi_N(x_{I,i}^W - \tau_h)$, where τ_h is chosen to match the pre-2023 mean share $\bar{h}$. Linearizing at $\bar{h}$ gives the model counterpart

$$\frac{\partial h_i}{\partial I_i} = \varphi\big[\Phi_N^{-1}(\bar{h})\big]\,\frac{\partial x_{I,i}^W}{\partial I_i},$$

where Φ_N is the standard-normal distribution function and φ is its density. The threshold rule provides the model counterparts for the pre-2023 sorting slope and the hiring response from 2023 onward. The ten targeted moments form $\hat{r}$, with model counterparts $r(\psi)$. The estimator solves

$$\hat{\psi} = \arg\min_{\psi:\, c\le 0} [\hat{r} - r(\psi)]' \widehat{\Omega}^{-1} [\hat{r} - r(\psi)],$$

where

$$\psi = (N_d, \zeta, \Delta p_A, w, \delta_g, \eta_L, c, \rho_I)'.$$

For coefficients estimated on overlapping firm samples, off-diagonal elements of $\widehat{\Omega}$ are constructed from stacked firm-level influence functions. Its diagonal contains the squared standard errors reported in Tables 4–6, with their stated clustering. Moments from disjoint samples have zero estimated covariance. Parameter uncertainty and prediction intervals come from 300 draws of the entire moment vector from its estimated joint normal distribution, with all eight parameters re-estimated at every draw.

To separate the cumulative 2023–2025 increase in sales per worker from subsequent accumulation, the propagation exercise begins from the model's 2025 allocation. The persistence parameter is 0.913, the autoregressive slope of year-demeaned log productivity in 2014–2022. The baseline assigns 30 percent of that persistence to learning from the matched worker, and Table A16 uses 15 and 45 percent as alternatives. The innovation covariance makes the pre-2023 firm distribution stationary, and the assignment map is recomputed after each annual update of firm capabilities.

Stability of the non-targeted sales-per-worker growth prediction. The post-2022 coefficient on information channel exposure in the sales-per-worker growth regression does not enter estimation. To assess whether any single targeted moment determines the model's

Table A16: Post-2025 Propagation through Learning from Matched Workers

Learning share	2023–25	+1 year	+3 years	+5 years	Long run
0.30	12.3	14.0	15.8	16.6	17.0
0.15	12.3	12.6	13.1	13.3	13.7
0.45	12.3	15.6	18.0	18.6	18.8

Notes. Entries are cumulative model-implied increases in sales per worker, in percentage points, relative to keeping signal precision and production efficiency at their pre-2023 values. Propagation begins from the 2025 allocation. The learning share is the fraction of the 2014–2022 productivity persistence assigned to the matched worker's contribution to next-period firm capabilities.

prediction for this coefficient, we re-estimate the model after omitting each target in turn. Table A17 reports the resulting predictions.

Table A17: Stability of the Sales-per-Worker Growth Prediction

Omitted target	Prediction	Omitted target	Prediction
Pre-2023 guidance-error intercept	1.520	Management-guidance error: Post × *F*-*K*-*I* composite	1.973
Guidance error–opacity slope	2.275	Management-guidance error: Post × opacity	2.170
Management-guidance error: Post ×*I*	2.462	Shrinkage error: Post ×*I*	2.242
Information-hire wage: Post ×*I*	1.719	Sales/worker: Post ×*P*	1.424
Employment: Post ×*P*	2.110	Pre-2023 slope: share of new hires in information roles	2.141

Notes. Target labels follow Panel B of Table 8. Each entry re-estimates the model after omitting the listed target and reports the resulting prediction for the *Post* × *I* sales-per-worker growth coefficient.

Structural parameter profiles. Table A18 shows how uncertainty about five of the estimated parameters affects the information-channel and reassignment shares.

Table A18: Structural Parameter Profiles

Parameter	95% profile values	Information-channel share (%)	Reassignment share (%)
c	[−2.02, −0.61]	46.3–49.3	14.6–39.4
w	[0.90, 2.31]	47.3–50.7	22.4–37.4
ζ	[0.29, 0.48]	45.2–54.0	30.1–32.4
Δp_A	[0.50, 50.00]	43.8–50.7	30.6–38.2
ρ_I	[1.34, 5.00]	48.7–49.2	26.0–38.8

Notes. For each row, we fix the listed parameter and re-estimate the remaining parameters. The reported values raise the re-estimated minimum-distance criterion by no more than 3.84, the 95 percent cutoff for one restriction. The estimation restricts $c \leq 0$, as described in Section 5.3. The upper endpoints shown for Δp_A and ρ_I are their parameter bounds. The information-channel share is the information-channel contribution divided by the combined increase. The reassignment share is the reassignment contribution divided by the information-channel contribution.

Fit among firms with the greatest exposure. A close fit to the average coefficients need not carry over to the most exposed firms. Table A19 therefore compares guidance, hiring, and sales-per-worker growth in the model and data for firms in the top quintile of information channel exposure.

Table A19: Top Quintile of Information Channel Exposure: Data and Model

Outcome	Data	Model	p-value: equal	Observations	Firms
Management-guidance error	-3.548^{***}	-4.603	0.356	5,041	869
	(0.840)	[-5.842, -3.236]			
Share of new hires in information roles	23.732^{***}	32.393	0.145	1,756	635
	(4.411)	[24.198, 37.971]			
Sales-per-worker growth	4.860^{***}	4.972	0.953	22,236	3,109
	(1.453)	[3.170, 7.100]			

Notes. Each row compares firms in the top quintile of information channel exposure with firms in the lower four quintiles. The data specification replaces $Post \times I$ with the top-quintile interaction, retains $Post \times P$, and otherwise follows the corresponding specification in Tables 4, 5, and 6. The model is evaluated on the exact firms in each regression. Data entries are coefficients with standard errors in parentheses. Model entries are predictions with 90 percent intervals from 300 joint draws in brackets. Guidance error, the share of new hires in information roles, and sales-per-worker growth are in percentage points. $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

Sensitivity of the channel decomposition. Table A20 varies the mapping from exposure to precision and production efficiency, the assignment calculation, the worker's payoff from a match, and the empirical moments used in estimation. It reports both the information-channel share and the part of that channel due to reassignment.

Baseline model extension: continuation value. In the baseline model, the old firm maximizes current profit and the young worker accepts the highest wage offered for her type, so her contribution to the firm's next-period capabilities in equation (4) does not affect assignment. In Jovanovic (2014), the young worker also values the expected profit of the firm she operates in the next period. The extension adds this expected profit, multiplied by a discount factor, to the worker's current wage. The resulting continuation value depends on the match, because the firm's next-period capabilities depend on both its current capabilities and the worker's skills. The last two rows of Table A20 re-estimate the model with discount factors of 0.50 and 0.95. The information-channel share is then 48.6 and 47.4 percent, compared with 49.2 percent in the baseline. With the continuation value, when learning from matched workers accounts for 15 or 45 percent, rather than 30 percent, of the 2014–

2022 year-to-year persistence in firms' sales per worker, the information-channel share ranges from 47.5 to 48.8 percent.

Table A20: Model Sensitivity

Specification	Combined	Information Channel	Information Channel Share (%)	Reassignment Share (%)
Panel A. Cross-firm heterogeneity				
No heterogeneity in initial precision	14.6	–	–	–
No heterogeneity in I	13.2	–	–	–
Baseline	12.3	6.0	49.2	31.3
Panel B. Alternative calibrations				
Log-precision intercept, $\bar{p}_A = -0.5$	16.5	10.5	63.6	29.5
Log-precision intercept, $\bar{p}_A = +0.5$	10.4	3.9	37.5	29.7
No heterogeneity in decision stakes, $\sigma_\ell = 0$	11.8	5.4	45.6	29.9
Labor share, $s_L = 0.55$	11.8	5.5	46.8	31.3
Labor share, $s_L = 0.65$	12.7	6.5	51.4	31.4
Guidance error on Post $\times I$, minus one SE	11.8	5.8	49.4	38.2
Guidance error on Post $\times I$, plus one SE	13.4	6.8	50.7	24.7
Uncorrelated worker skills	12.6	6.5	51.2	31.1
Production channel exposure floor at minimum	17.1	6.0	35.4	31.3
Production channel exposure floor at twenty-fifth percentile	11.1	6.0	54.4	31.3
Independent forecast-error variance, 25%	12.7	6.2	48.7	31.5
Independent forecast-error variance, 50%	9.4	4.1	44.0	1.4
Panel C. Assignment and alternative moments				
No cross-task interaction, $c = 0$	7.8	3.1	39.8	0.8
Mean interaction matrix positive definite	9.6	4.4	46.0	1.4
All initial interaction matrices positive definite	7.8	3.1	40.2	0.8
Discrete optimal assignment	12.1	5.9	48.6	29.7
Production channel exposure floor excluded from assignment	12.2	6.0	49.2	31.5
Sales-per-worker growth moment included	12.3	6.4	52.0	29.3
H-1B wages by task of hire, c unrestricted	12.6	5.6	44.4	26.5
Benchmark forecast: 25% of information-channel improvement	12.4	6.3	50.9	31.8
Benchmark forecast: 50% of information-channel improvement	12.2	6.3	51.3	31.8
Shrinkage-benchmark moment excluded	12.4	6.3	50.6	31.7
Worker's discount factor for next-period firm profit, 0.50	12.3	6.0	48.6	30.0
Worker's discount factor for next-period firm profit, 0.95	12.3	5.9	47.4	26.9

Notes. The table reports model-implied cumulative 2023–2025 differences in sales-per-worker levels for the sample, in percentage points. The information-channel share is the information-channel contribution divided by the combined increase. The reassignment share is the reassignment contribution divided by the information-channel contribution. Panels B and C vary one input or moment at a time and re-estimate the model. The discrete-assignment row instead holds the estimated parameters fixed and replaces the assignment map in Theorem 1 with the average of three discrete assignments on the observed firm sample and the same simulated worker distribution. The two benchmark rows allow the stated share of the information-channel improvement to enter the shrinkage benchmark forecast. The row for H-1B wages by task of hire uses all four coefficients in Panel C of Table A13 as estimated moments and removes the sign restriction on c, which gives $c = -1.01$. The last two rows add the worker's expected profit from operating the firm next period, multiplied by the stated discount factor, to her current wage. The covariance of firm capabilities is held at its observed pre-2023 value. As in the baseline, learning from matched workers accounts for 30 percent of the 2014–2022 year-to-year persistence in firms' sales per worker.