Hedging AI Shocks to Human Capital

Xijin Huang, Jose Vicente Martinez, Michael Pokojovy, and Matthias Qian*

Preliminary draft. Please do not cite or circulate without the authors' permission.

ABSTRACT

We study whether equity markets can hedge human capital risk arising from skill-specific shocks. Using job postings data from 2013–2023, we estimate shadow wages for more than 500 skills and identify innovations in skill demand and compensation. We construct equity-based projection portfolios, including an AI portfolio based on firms' AI patent activity, and test their ability to hedge these shocks. Both in-sample and out-of-sample results indicate economically meaningful hedging performance for a non-trivial subset of skills. Embedding these portfolios in a life-cycle model and solving the dynamic consumption–investment problem, we find that welfare gains for exposed individuals under the estimated skill processes are modest. However, when allowing for a moderate probability of skill crashes, such as those potentially induced by AI, these portfolios become substantially more valuable.

March 2026

Keywords: Human capital and labor risk, Skill value shocks, Hedging portfolios, Technological change, Artificial intelligence

JEL classification: G11, J24, J30, O33

* Xijin Huang (xijin.huang@uconn.edu) and Jose Vicente Martinez (jose.v.martinez@uconn.edu) are with the University of Connecticut, School of Business, Michael Pokojovy (mpokojovy@odu.edu) is with Old Dominion University, School of Data Science, and Matthias Qian (matthias.qian@sbs.ox.ac.uk) is with ESMT Berlin. We are grateful to Anna Chevychalova for research assistance.

I INTRODUCTION

Human capital, understood as a collection of skills, knowledge, health, and values that contribute to making people productive (e.g., Becker, 1964; Benzoni & Chyruk, 2015), is the largest component of most people's portfolios. Empirical estimates suggest that it accounts for roughly 60% to more than 90% of total wealth for the representative investor (Jorgenson & Fraumeni, 1989; Palacios, 2015). Yet human capital is also inherently risky. Its value is exposed to shocks arising from technological change and shifting skill demands—shocks that, with the rapid diffusion of artificial intelligence (AI), may unfold quickly and unevenly across the workforce and are difficult to insure through traditional labor market institutions.[1]

In this study, we investigate whether equity markets can hedge human capital risk. We depart from the standard approach that treats labor income as a single homogeneous process and instead adopt a granular view of human capital as a portfolio of skills, each generating its own stochastic income stream. This perspective reflects both the nature of human capital as a set of endowed or acquired skills and a growing consensus in the labor economics literature that skills, rather than occupations or industries, are the fundamental units through which technological change affects workers (Autor et al., 2003; Acemoglu & Autor, 2011; Deming, 2017). It also naturally implies heterogeneity in exposure to technological shocks: some skills are complemented by AI, others are substituted, and many are largely unaffected. We therefore test whether equity portfolios, particularly portfolios constructed to load on AI- and technology-related risks and opportunities, can hedge innovations to skill-level labor income.

[1] According to Frey & Osborne (2017), 47 percent of jobs in the US economy are at risk of being automated by advances in AI-related fields. Eloundou et al. (2023) argue that around 80% of the U.S. workforce could have at least 10% of their work tasks affected by the introduction of LLMs, while approximately 19% of workers may see at least 50% of their tasks impacted and up to 56% of all worker tasks in the US could be completed significantly faster at the same level of quality as a result of the adoption of LLM and complementary technologies.

Our approach builds on and extends the mimicking portfolio framework of Lamont (2001) and Engle et al. (2020), applying it to the labor market at the level of individual skills and occupations. We combine large-scale job postings data with asset pricing techniques to construct hedge portfolios designed to offset innovations in skill remuneration, skill demand, and related measures of labor market value. In doing so, we bridge two literatures that have largely evolved separately: the study of human capital risk in labor economics and the study of hedging non-traded risks using financial markets in asset pricing.

Methodologically, we proceed in three steps. First, we construct monthly time series of skill-level labor market outcomes using job postings data from Lightcast covering the United States from 2013 to 2023. Using hedonic wage regressions in the spirit of Rosen (1974) and Tinbergen (1956), we estimate monthly "shadow wages" for over 500 skills, controlling for geography and industry composition. From these series, as well as from skill posting counts and skill value measures combining skill postings and wages, we extract innovations using ARIMA and SARIMA models. These innovations capture unexpected changes in the remuneration and demand for specific skills and serve as our hedge targets.

Second, we construct a set of equity-based projection portfolios designed to capture exposure to technological and broader economic risks. Our primary projection portfolio is an AI portfolio formed by sorting firms on their intensity of AI-related patenting, scaled by market capitalization. We complement this portfolio with market, size, and value projection portfolios constructed following Engle et al. (2020). Using the mimicking portfolio approach of Lamont (2001) and the hedging methodology proposed by Engle et al. (2020), we regress skill-level innovations on lagged projection portfolio returns, allowing for delayed adjustment in labor

markets. The coefficients from this regression determine the weights of hedge portfolios designed to offset skill-specific shocks.

Third, we evaluate the performance of these hedge portfolios both in sample and out of sample. In-sample tests assess how much variation in skill innovations can be hedged using contemporaneously estimated portfolio weights. Out-of-sample tests provide a more stringent assessment by estimating hedge portfolios using rolling one-year and three-year training windows and examining their ability to hedge subsequent skill innovations. Finally, we embed the estimated hedge portfolios into a portfolio choice framework with CRRA utility to quantify their economic value for individuals with different skill profiles, different risk tolerances, and at different points in their life cycle.

Our results yield several key findings. First, in-sample regressions indicate that equity portfolios sorted on AI exposure meaningfully hedge skill-level labor market shocks for a nontrivial subset of skills. On average, the projection portfolios explain around 10–12% of the variation in skill wage and skill demand innovations, with substantially higher explanatory power for skills more exposed to technological change.[2] The AI projection portfolio plays a dominant role in these results: its contribution is statistically significant for a large fraction of skills, far exceeding what would be expected by chance, while traditional market, size, and value portfolios contribute relatively little. These findings suggest that AI-related equity exposures capture economically relevant variation in labor market outcomes.

Second, out-of-sample performance remains economically and statistically meaningful. Hedge portfolios constructed using rolling training windows exhibit positive correlations with three-month-ahead skill innovations across a wide range of hedge targets, including skill wages,

[2] These magnitudes are comparable to those reported by Engle et al. (2020) for climate risk hedging.

posting counts, and value measures. The fraction of skills for which these correlations are statistically significant substantially exceeds random benchmarks, particularly for skill postings and wage-based measures. Performance is stronger when using shorter training windows, consistent with time variation in the relationship between equity markets and labor market outcomes. Notably, these results likely understate the true effectiveness of hedging, as many skills are not expected a priori to be exposed to the shocks captured by our projection portfolios.

Third, the economic value of these hedge portfolios depends on the nature of skill risk. Embedding the estimated skill processes in a life-cycle model and solving the dynamic consumption–investment problem, we find that welfare gains from access to hedge portfolios—measured as increases in certainty-equivalent consumption—are modest. This reflects the relatively low volatility of estimated log skill income innovations and the ability of individuals to smooth consumption using risk-free assets and the market portfolio; even a perfect hedge against skill innovations would deliver limited additional benefits in this scenario. However, when we allow for a moderate probability of large skill declines, such as those potentially induced by technological disruptions like AI, the value of the hedge portfolios increases substantially, as they provide protection against otherwise difficult-to-insure downside risk.

Overall, our findings suggest that financial markets—particularly portfolios capturing exposure to AI-related technologies—can partially hedge human capital risk at the level of individual skills. While such hedging is necessarily imperfect and heterogeneous, it nonetheless represents a potentially important mechanism through which individuals can insure against labor market disruptions induced by technological change.

Our analysis builds on and extends several strands of prior literature. First, it relates to work on hedging non-traded risks using financial markets. Engle et al. (2020) propose constructing

mimicking portfolios (Lamont, 2001) to hedge climate risk. We extend this framework by applying it to labor market outcomes at the level of individual skills and occupations. Focusing on skills rather than aggregate wages or industry-level income allows us to capture the heterogeneous effects of technological change across skills (Eloundou et al., 2023; Hampole et al., 2025), and distinguishes our contribution from earlier studies that concentrate on aggregate wage streams (Davis & Willen, 2000).

Second, our paper contributes to the growing literature on AI, technological change, and labor markets. A large body of work documents how automation and AI reshape employment, task composition, and wages (e.g., Autor et al., 2015; Acemoglu & Restrepo, 2020; Babina et al., 2023, 2024). Kogan et al. (2021) show that workers exposed to innovative firms experience higher income volatility, while Kogan et al. (2023) link worker displacement to firm-level AI adoption and patenting activity. We complement this literature by examining whether the risks associated with these technological changes can be hedged using traded financial assets. Our use of skill-level wage and demand innovations offers a higher-resolution lens on heterogeneity in technological exposure and hedgeability.

Third, our work speaks to the literature on portfolio choice in the presence of labor income risk. A long tradition models labor income as a non-traded asset that affects optimal portfolio allocations (Heaton & Lucas, 2000; Viceira, 2001; Cocco, et al., 2005; and more recently, Duarte et al., 2021; Catherine, 2022). Whereas this literature emphasizes diversification between financial wealth and labor income, our study focuses on portfolios explicitly constructed to hedge human capital and labor income risk. Our results suggest that at least part of skill-level labor income risk, particularly risk related to AI, may be hedgeable using traded assets, with meaningful implications for welfare.

II. IDEA DEVELOPMENT AND METHODOLOGY

A. Human Capital as a Portfolio of Skills

We conceptualize human capital as a portfolio of skills, each generating a stochastic income stream. This perspective aligns with standard definitions of human capital as a collection of productive skills (Becker, 1964; Benzoni & Chyruk, 2015) and is particularly appropriate in our setting, given evidence that skills—rather than occupations or industries—constitute the primary channels through which technological change affects workers (Autor, Levy, & Murnane, 2003; Acemoglu & Autor, 2011; Deming, 2017). For tractability, and in line with much of the portfolio choice literature with labor income risk (Heaton & Lucas, 2000; Viceira, 2001; Cocco, Gomes, & Maenhout, 2005), we treat skill income as exogenous from the individual's perspective and abstract from endogenous skill acquisition, retraining, and labor supply decisions.

We operationalize this framework by estimating skill shadow wages, defined as the implicit market remuneration of individual skills. Using large-scale job postings data, we estimate monthly hedonic wage regressions in the spirit of Tinbergen (1956) and Rosen (1974). From the resulting time series, we extract innovations to skill shadow wages by fitting autoregressive moving-average models (Davis & Willen, 2014). We focus on innovations because they capture the component of skill-level remuneration that is unanticipated given past information and therefore the one relevant for risk and hedging.

B. Hedging Skill Risk using Mimicking Portfolios

Our objective is to assess whether equity markets can hedge shocks to skill-level human capital. To do so, we adopt the mimicking portfolio framework of Lamont (2001), as implemented by Engle et al. (2020) in the context of climate risk. The central idea is to construct portfolios of traded

assets whose returns covary with innovations to a non-traded income stream, in our case, skill shadow wages.

For a given skill, a hedge portfolio is designed to deliver positive returns when adverse innovations to that skill occur, thereby partially offsetting losses in human capital. An investor who holds such a portfolio can reduce the variance of lifetime income and improve welfare, even if the hedge portfolio itself earns no abnormal expected return. Our analysis therefore focuses on hedgeability, not on whether skill risk is priced or commands a risk premium in equilibrium.

A key empirical challenge is dimensionality. The cross-section of tradable assets is large relative to the available time-series observations on skill innovations. Estimating unrestricted portfolio weights would therefore be infeasible and prone to severe overfitting. To address this issue, we follow Engle et al. (2020) and restrict portfolio weights to be functions of firm characteristics that proxy for exposure to underlying economic forces relevant for labor markets. This approach yields parsimonious, economically interpretable hedge portfolios that are feasible to estimate in finite samples.

C. Projection Portfolios and AI-Related Technological Exposure

Our primary focus is on AI-related technological exposure, motivated by growing evidence that AI generates large, uneven, and potentially rapid shocks to labor demand and wages across skills (Frey & Osborne, 2017; Acemoglu & Restrepo, 2020; Eloundou et al., 2023; Mäkelä & Stephany, 2024; Hampole et al., 2025). Rather than attempting to identify pure AI shocks, we ask a more modest and empirically tractable question: Do equity portfolios capturing exposure to AI-related technologies hedge skill-level labor market risk?

To capture AI-related exposure in equity markets, we construct a novel AI projection portfolio based on firm-level AI patenting activity.[3] Firms are weighted by their cumulative stock of AI patents, scaled by market capitalization and demeaned cross-sectionally to produce a zero-investment portfolio. This portfolio is not intended to perfectly span AI-related technological shocks. AI patenting is a noisy and imperfect proxy, reflecting differences in patent quality, commercialization, and timing. Moreover, technological progress does not necessarily translate one-for-one into firm value: gains may be competed away, passed through to consumers, or offset by creative destruction (Kogan, Papanikolaou, & Stoffman, 2020). For these reasons, whether AI-loaded equities provide an effective hedge for skill-level labor market risk is ultimately an empirical question.

To complement the AI projection portfolio, we also include market, size, and value projection portfolios, constructed following Engle et al. (2020). These portfolios capture broader sources of systematic risk and allow us to assess whether AI-related exposure provides incremental hedging power beyond standard asset pricing factors.

D. Estimation and Evaluation of Hedge Portfolios

We estimate hedge portfolios by regressing innovations to skill-level labor market outcomes on lagged returns of the projection portfolios. Allowing for a lag reflects the delayed adjustment of labor markets to new information.

We evaluate hedging performance along several dimensions. In-sample analyses assess whether projection portfolios explain economically meaningful variation in skill innovations. While informative, these results may overstate true hedging potential due to estimation error. We therefore place particular emphasis on out-of-sample performance, estimating hedge portfolio

[3] Patent-based measures of innovation have a long tradition in finance as proxies for technological exposure and innovative capacity (Hirshleifer, Hsu, & Li, 2013; Kogan et al., 2017).

weights using rolling training windows and evaluating their ability to hedge subsequent skill innovations. This approach mirrors the real-time problem faced by investors.

Finally, we embed the estimated hedge portfolios into a dynamic portfolio choice problem to assess their economic value.[4] A hedge portfolio that reduces labor income risk may not be that attractive if it introduces substantial financial risk. To quantify the value of these hedge portfolios, we solve a lifecycle consumption–portfolio problem with non-traded labor income risk and CRRA preferences. We quantify welfare effects using certainty-equivalent consumption and wealth measures and explore sensitivity to assumptions about expected hedge portfolio returns and income risk.

III. SKILLS, OCCUPATIONS, WAGES, AND POSTINGS

We use job postings data from Lightcast to construct a monthly time series (2013–2023) of skill "shadow" wages and postings.[5] Compared with résumé data, job postings provide much greater detail: résumés typically list only job titles and rarely include sufficiently detailed descriptions of work activities or required skills. Our analysis relies heavily on the skills extracted by Lightcast from the descriptions of tasks contained in these job postings. All wage measures also come directly from Lightcast job postings. These postings typically report a salary range, and we use the midpoint of the range. To ensure comparability over time, we adjust advertised salaries for inflation using the Personal Consumption Expenditures (PCE) index.

There are, of course, some caveats to using this data to infer skills' remuneration and usage. Posted vacancies do not always translate into hires, and some hires occur without public postings

[4] To solve the resulting high-dimensional and nonlinear optimization problem, we rely on recent advances in deep learning–based dynamic programming (Kopeliovich & Pokojovy, 2024). The use of machine learning here is purely computational.

[5] See Babina et al. (2024) and Chen, Wu, and Yang (2025) for examples of the use of this dataset (previously known as Burning Glass) in finance.

(e.g., via recruiting firms). Nonetheless, postings and hires are generally highly correlated, making postings a reasonable proxy for labor demand and a useful source of information for spot wages.

A. Estimation of Skill Shadow Wages

We obtain skill shadow wages from a regression of real job posting salaries or log salaries on a full set of skill indicator variables, i.e., a hedonic wage function (see Rosen, 1974).[6] A separate regression is estimated for each month. In addition to skill indicator variables, these regressions also include a full set of state and industry indicators.

Formally, the skill shadow wages model we estimate for each month *t* is as follows:

$$w_{it} = \omega_t' s_{it} + \gamma_t' x_{it} + \varepsilon_{it}, \qquad t = 1,2,\dots,T \quad (1)$$

where w_{it} is the advertised salary or log salary in job posting *i* published on month *t* adjusted by the PCE index (we use a July 2018=1 base series, as a result, salaries are expressed in constant 2018 purchasing-power-equivalent dollars), s_{it} is a J × 1 vector of skill indicator variables such that each indicator takes the value one if job posting *i* references the skill and zero otherwise, ω_t is a J × 1 vector of month *t* skill coefficients, x_{it} is a K × 1 vector of control variables (full sets of state and industry indicators), γ_t is a K × 1 vector of control variable coefficients, and ε_{it} is an idiosyncratic deviation.

Our data contains several million job postings (78,412,578 between 2013 and 2023) and thousands of different skills (16,109 in the same time span). Skills are not mutually exclusive, and a job posting may reference or require multiple skills. For computational reasons, we focus on the 525 most frequent skills, including an indicator variable for each one in our regressions, and

[6] In economics (especially labor economics and applied microeconometrics), it is much more common to model log wages (log wages) rather than raw wages. This is because the log transformation helps deal with skewness, mitigates heteroskedasticity, captures proportional rather than absolute effects, and allows coefficients to be interpreted as percentage changes in earnings.

summarize the remaining skills using a count variable ("other skills count"). The model is estimated using ordinary least squares (OLS).

In model (1), the coefficient associated with each skill indicator variable reflects the average (log) remuneration of that skill in the omitted state and industry, in the context of a (log) linear model. We define the "skill shadow wage" for each month in the USA and across industries as the sum of the coefficient of the skill indicator variable, the weighted average of the state coefficients (including zero for the omitted state), and the weighted average of the industry coefficients (including zero for the omitted industry). In these specifications, state coefficients (including zero for the omitted state) are weighted by state job postings in the overall sample (or population) and industry coefficients are weighted in a similar way. Importantly, industry and state weights are kept constant throughout the entire sample to avoid changes in the geographic or industry composition of the labor force contaminating the skill wage measures (and, later on, the innovations to these measures).

Table 1 presents the ten skills with the highest and lowest average wages among the 100 most frequent skills in our sample. For each skill, the table reports its name, total number of appearances in the sample, the percentage of total skill postings, the average wage during our sample period, and the skill classification. Skills are classified into one of five categories (analytical, social, managerial, mechanical, and fundamental skills) using ChatGPT and following Kochhar (2023).

The skills with the highest implied wages are registered nurse, CDL class A license, truck driver, negotiation, and project management. And those with the lowest implied wages are merchandising, warehousing, retail sales, restaurant operation, and housekeeping.[7]

[7] A small number of skills show negative real shadow wages in certain months, which may stem from estimation errors or model misspecifications. However, none of the 100 most common skills in the sample exhibit negative average

B. Innovations to Skill Shadow Wages and Postings

In our analysis, we focus on the 100 most common skills. We extract innovations to these skills by fitting ARIMA models to monthly skill measures (see Davis and Willen, 2014). For real skill shadow wages and real skill log wages, we follow the micro-earnings dynamics literature (Baker, 1997; Guvenen, 2009) and model these series as potentially persistent but stationary processes.[8] As shown in Table 2, the ARIMA(1,0,1) specification most frequently provides the best fit across skills. Accordingly, we adopt an ARIMA(1,0,1) model for all skill time series.

Formally, for each skill shadow wage *j,* we estimate the following model:

$$\omega_{jt} = d_j + \phi_j \omega_{jt-1} + \iota_{jt} + \theta_j \iota_{jt-1} \qquad j = 1,2,\dots,J \qquad (2)$$

where d_j is the drift term, ϕ_j is the AR coefficient, θ_j is the MA coefficient, and ι_{jt} is the time *t* skill *j* shadow wage innovation.

Our primary interest is in estimating real skill shadow wage or log wage innovations. However, we also explore innovations to skill quantities (numbers of skill postings), values (number of skill postings times wages) and log values, because in labor markets, price adjustments often follow quantity adjustments, particularly when wages are sticky (see Stiglitz, 1984).[9]

For this purpose, we work with skill counts adjusted by population growth. Adjusted skill counts are obtained by dividing monthly skill counts (in job postings) by the monthly working age population. For this purpose, we use estimates of the working age population from the Organization for Economic Co-operation and Development, Infra-Annual Labor Statistics:

shadow wages over the sample period, nor do they record negative shadow remunerations in more than half of the observations.

[8] Our tests in Table 2 indicate that while some skill shadow wage and log wage series are stationary, for others the Augmented Dickey–Fuller (ADF) test fails to reject the null of a unit root. It is well known, however, that unit root tests have low power against highly persistent but stationary alternatives.

[9] See also Akerlof et al. (1996), Kahn (1997), Bewley (1999), and Grigsby et al (2021), among others, for arguments and evidence of downward wage rigidity.

Working-Age Population Total: From 15 to 64 Years for United States (retrieved from FRED, Federal Reserve Bank of St. Louis). Using these estimates, we construct an index by dividing each monthly value of the series by the value in July 2018, thereby creating a July 2018=1 base series. In this way, we capture fluctuations in skill employment (or demand) over time adjusted by population growth.

Our analysis suggests that the number of job postings per skill is non-stationary (for all skills) and displays clear seasonality, with the SARIMA(1,1,0) × (0,1,1,12) specification most often yielding the best fit (see Table 2). For this reason, we fit a SARIMA(1,1,0)×(0,1,1,12) to all adjusted skill count time series to obtain innovations to the number of job postings mentioning each skill in the entire USA and across industries.

The skill value (wages multiplied by quantities) and log value series show similar behavior. Therefore, we also fit SARIMA(1,1,0) × (0,1,1,12) models to all adjusted skill value and log value time series to obtain the corresponding innovations in skill (log) values.

C. Occupation Wages, Postings, and their Innovations

In addition to analyzing skill shadow wages, skill posting counts, and values, we also examine occupation wages, occupation posting counts, and values. While skills represent more stable, long-term attributes of individuals and therefore have a stronger theoretical justification, estimates of occupation wages, occupation postings, and values tend to be more precise and less model-dependent, making them a useful validation exercise for our analysis.

As with the skill-level analysis, we consider both raw and log-transformed wages. Specifically, we estimate a model similar to equation (1), where the wage (or log wage) associated with each posting is regressed on occupation indicators and a full set of state and industry

indicators. The main difference from the skill specifications is that occupations, unlike skills, are mutually exclusive.

Table 3 reports the ten occupations with the highest and lowest average wages. Among the 100 most common jobs, those with the highest average wages include nurse practitioners, computer systems engineers, software developers, treasurers and controllers, and critical care nurses. In contrast, the lowest-paying jobs include cashiers, dishwashers, fast food and counter workers, hostesses and hosts at restaurants, lounges, and coffee shops, as well as waiters and waitresses.

Similar to the skill wage series, most occupation wage (and log wage) time series are stationary, with an ARIMA(1,0,1) model providing the best fit in most cases. Accordingly, we use an ARIMA(1,0,1) specification to extract innovations for all occupation wage time series.

Our analysis also indicates that the number of job postings per occupation, after adjusting for population growth, is typically non-stationary and exhibits clear seasonal patterns. The SARIMA(1,1,0) × (0,1,1,12) model most frequently offers the best fit (see Table 4). Therefore, we apply this specification to all adjusted job posting count time series to obtain the corresponding innovations.

Finally, the occupation value and log value series show similar behavior. Therefore, we also fit SARIMA(1,1,0) × (0,1,1,12) models to all adjusted occupation value and log value time series to obtain the corresponding innovations in occupation (log) values.

IV. FORMATION AND ANALYSIS OF HEDGE PORTFOLIOS

The main goal of our study is to examine whether equity markets can serve as a hedge against shocks to human capital. Specifically, we analyze how a set of well-diversified projection portfolios, each capturing different dimensions of risk, perform in this role. We place particular

focus on risks to human capital arising from AI and new technologies, as well as on portfolios specifically designed to hedge against such risks, i.e., those expected to exhibit high correlation with these shocks. We also emphasize the heterogeneous nature of human capital, stemming from differences in skills, knowledge, and abilities, which generates dispersion not only in income but also in exposure to AI and related risks.

Given our focus on AI-related risk, our analysis begins in July 2018. This period marks a turning point in how businesses and the public perceived AI. As shown in Figure 1, the frequency of Wall Street Journal articles referencing "artificial intelligence," "AI," "machine learning," "ML," and "deep learning" increased sharply, from occasional mentions to regular coverage, around this time. This surge coincided with the White House's first "Artificial Intelligence for American Industry" summit in May 2018, an event that signaled top-level policy attention to AI, and the subsequent release of GPT-1 in June 2018.

A. Hedge Targets and Projection Portfolios

Our hedge targets alternatively include innovations to different skill shadow wages, occupation wages, skill posting counts, occupation posting counts, skill values (defined as wages multiplied by quantities), and occupation values. We define each hedge target in both real dollar terms and real log-dollar terms. Although dollar-based hedge targets may appear more natural in finance applications, any portfolio that hedges log-dollar wages and their innovations will also hedge their dollar counterparts. This approach aligns with the labor economics literature, which emphasizes that log wages exhibit more stable and well-behaved statistical properties than raw wages.

To construct hedge portfolios, we use U.S. equities as the underlying assets. Monthly individual stock return data are obtained from CRSP, including only common equity securities (share codes 10 and 11) listed on the NYSE, AMEX, and NASDAQ. Following Engle et al. (2020),

we exclude penny stocks, those priced below $5 at the time of portfolio formation, and microcap stocks, defined as firms within the bottom 20% of NYSE market capitalization. Nominal returns are converted to real terms by deflating with the PCE index.

Using these assets, we build four projection portfolios. The first is an AI projection portfolio, formed by weighting stocks according to their AI capital, proxied by the number of AI patents per firm. For this, we use the AI patent dataset constructed by Pairolero et al. (2025), matched to publicly traded firms via the bridge file developed by Kogan et al. (2017). Each firm's cumulative patent count (its patent stock) is scaled by market capitalization. We obtain the vector of portfolio weights for the AI projection portfolio, z_t^{AI}, by cross-sectionally demeaning each firm's monthly patent stock. The resulting portfolio is a zero-investment portfolio.

On average, the largest average portfolio position by 2-digit SIC code industry classification in this projection portfolio is "Chemicals and Allied Products," with "Electronic, Electrical Equipment and Components," "Measuring, Analyzing, and Controlling Instruments; Photography, Medical and Optical Goods; Watches and Clocks", and Industrial and Commercial Machinery and Computer Equipment" also in the top 5. By contrast, the largest short positions include "Electric, Gas, and Sanitary Services," "Business Services," and "Insurance Carriers," in what is a fairly industry-diversified portfolio (see Appendix Table 1).

In addition to the AI projection portfolio, we also construct market, $z_{t-1}^{MKT\prime}$, value, z_{t-1}^{HML}, and size, z_{t-1}^{SI}, portfolios following the procedure outlined by Engle et al. (2020). We label these EGKLS-style Fama–French projection portfolios MKT, HML, and SMB, respectively.

Our portfolios are not necessarily intended to be the definitive projection portfolios; refinements such as incorporating alternative AI-related variables or improved patent-based measures (see Hirshleifer, Hsu, and Li, 2013) could yield better performance. Nonetheless, they

should capture a reasonable representation of the cross-section of risk, particularly the risks currently impacting human capital.

B. Forming Hedge Portfolios

We build portfolios designed to hedge shocks to human capital using the mimicking portfolio approach advocated by Lamont (2001). To do this, we regress each hedge target on our set of projection portfolios, as in Engle et al. (2020). Specifically, for each skill *j* shadow wage innovation, occupation *j* wage innovation, or skill/occupation *j* count or value, we estimate the following model:

$$\iota_{jt+k} = \xi_j + \pi_j^{AI} z_{t-1}^{AI\prime} r_t + \pi_j^{SIZE} z_{t-1}^{SIZE\prime} r_t + \pi_j^{HML} z_{t-1}^{HML\prime} r_t + \pi_j^{MKT} z_{t-1}^{MKT\prime} r_t + e_{jt} \qquad (3)$$

where r_t is an N × 1 vector of time *t* excess stock returns over the risk-free rate, z_{t-1}^{AI}, z_{t-1}^{SIZE}, z_{t-1}^{HML}, and z_{t-1}^{MKT} are N × 1 vectors of AI, SIZE, HML, and MKT firm characteristics (as of time *t*-1), and π_j^{AI}, π_j^{SIZE}, π_j^{HML}, and π_j^{MKT} are scalars that capture the weights of each projection portfolio in the mimicking (hedge) portfolio for ι_{jt}.

In these regressions, we incorporate a lag of k=3 periods between stock returns and human capital innovations to reflect the delayed adjustment of labor markets to new information (i.e., current returns hedge innovations k=3 months ahead). All models are estimated using monthly data and ordinary least squares (OLS).[10]

C. In-Sample Results

We begin by examining whether AI-, market-, size-, and value-sorted portfolios help hedge human capital innovations in-sample. Table 5 presents summary statistics from regressions of skill shadow wages, skill posting counts, and skill values on AI and EGKLS-style Fama–French hedge

[10] We tested several lag lengths and ranges. The results presented in the paper are based on a 3-month lag between stock market and labor market outcomes. Comparable levels of significance are found when using lag ranges of 2–4 (i.e., projection portfolio returns lagged 2, 3, and 4 months in the same regression), 3–4, or 3–5 months.

portfolios. Results in this table indicate that, for skill (shadow) wage and log wage models, average R-squares stand at 0.12 for the 2018-2023 period. These are average R-squares calculated across the 100 most common skills in the Lighcast dataset. These results suggest that using these hedge portfolios, the average individual should be able to hedge 10% of the in-sample variation in wages or log wages (some more, some less).[11] Although these R-squared values and implied hedging abilities may appear modest, they are averages: individuals whose skills are more exposed to the factors captured by our projection portfolios should be able to hedge more of their human capital shocks, while others will hedge predictably less.

An analysis of the contributions of each projection portfolio indicates that the AI portfolio is the one that most meaningfully contributes to hedging wage and log wage risk. The AI portfolio is significant ($p < 0.10$) in 32% of skill wage regressions, ($p < 0.05$) in 26%, and ($p < 0.01$) in 19%. The significance of the AI portfolio is even more pronounced in the log wage regressions, with 42%, 36%, and 20% of the analyzed skills exhibiting a significant contribution of this portfolio at the 10%, 5%, and 1% levels, respectively. These numbers are far greater than expected by chance, as confirmed by binomial probability tests.[12] This is especially notable given that not all skills are expected a priori to be exposed to AI shocks. By contrast, the contributions of the MKT, HML, and SMB portfolios are generally statistically indistinguishable from zero.

When the hedge target is innovations in skill posting counts, the average R-square is 0.10, slightly below that of the skill wage regressions, but with a comparable proportion of skills for which the AI portfolio is statistically significant. Similar patterns emerge when we use skill value and log value indices, defined as the product (or log product) of skill wages and posting quantities,

[11] For comparison, the R-square in Engle et al (2020)'s best climate hedge regressions is in the 0.15—0.19 range.
[12] The reported binomial probability (test) indicates the probability of observing the number of observed significant values or an even more extreme number, under the null hypothesis of no relationship between projection portfolios and hedge targets, given the specified p-value thresholds.

as hedge targets. Using occupation-based hedge targets, that is, occupation wages or log wages, occupation posting counts, or occupation value measures, yield similar results (see Table 6).

Overall, these results suggest that AI sorted portfolios can partially hedge shocks, possibly AI shocks, to at least some skills. Table 7 lists the skills for which such hedging appears most effective (i.e., those skills most negatively exposed to the AI portfolio). These include spreadsheets, telephone skills, multilingualism, and data entry, along with less intuitively AI-linked skills such as retail sales and enthusiasm.

Similarly, Table 8 identifies occupations that benefit most from AI hedging during out sample period, including intuitive candidates such as clerks and receptionists, as well as less expected ones such as retail salespersons—potentially influenced by the inclusion of the COVID-19 period in our sample. It is important to notice that the AI projection portfolio used in this study is not a generative AI projection portfolio but rather a general AI projection portfolio. Finally, Tables 7 and 8 indicate that none of the most frequent skills or occupation values are reliably positively associated with AI portfolio returns, implying that no individual equipped with those skills would benefit from shorting the AI portfolio.

D. Out-of-Sample Results

Arguably, the most meaningful test of a hedge portfolio's performance is its ability to hedge out of sample, that is, to hedge innovations in periods not used to estimate the portfolio weights. To assess the out-of-sample performance of our hedging strategy, for each period t we estimate regression (3) using data from either t–13 to t–1 (a one-year training window) or from t–37 to t–1 (a three-year training window). Based on these estimates, we construct the hedge portfolio by combining the AI projection portfolio and the three EGKLS-style Fama–French projection

portfolios, each weighted by its estimated coefficient. We then examine the correlation between the returns of this hedge portfolio and the innovation measures for each skill.

In our analysis, projection portfolios are lagged by three months, which means that the hedge portfolio's month t return is designed to hedge three-month-ahead skill innovations or innovations that will impact labor markets three months after. Thus, in our out-of-sample evaluation, we compare the hedge portfolio return in month t with the skill innovations in month t+3.

The average correlation between skill-level innovations and out-of-sample hedge portfolio returns (see Table 9, Panel A) ranges from 0.05 for skill log value targets to 0.20 for skill posting count targets.[13] The percentage of skills for which the hedge portfolio exhibits a positive and significant correlation with three-month-ahead innovations varies between 6% and 53% ($p < 0.10$), 2% and 34% ($p < 0.05$), and 1% and 11% ($p < 0.01$). Most hedge targets lie in the middle to upper portions of these ranges, whereas skill log value targets display the weakest performance. Hedging performance is generally stronger when shorter (12-month) training windows are used. These results usually far exceed what would be expected by chance, as confirmed by binomial probability tests. This is especially the case for skill postings, skill shadow wage, and log shadow wage targets. As noted in the in-sample analysis, many skills are not expected a priori to be exposed to the shocks captured by our projection portfolios, so reported significance percentages likely understate the true significance of these results.

When restricting the analysis to skills that were significant in-sample (at the 10% level), both the average correlations and the fraction of skills effectively hedged out of sample increase (see Table 9, Panel B). Although this exercise introduces some look-ahead bias in skill selection

[13] For comparison, Engle et al. (2020) report correlations between climate news innovations (targets) and out-of-sample hedge portfolio returns ranging from 0.17 to 0.20 for their best-performing models.

and therefore does not constitute a valid hedging strategy, it remains informative and provides context for the more detailed portfolio analysis in the next section.

Using occupation-based hedge targets, such as job wages, log wages, posting counts, or job value measures, produces qualitatively similar but generally weaker results (see Table 10). The average correlation between innovations and lagged hedge portfolio returns ranges from 0.04 to 0.13, depending on the target, and is considerably higher when focusing on occupations that performed well in-sample. Overall, these findings suggest that hedge portfolios can partially hedge real-time shocks to human capital, including AI-driven shocks, for at least a subset of skills or occupations.

V. VALUE OF THE HEDGE PORTFOLIO TO INDIVIDUAL INVESTORS

In light of the results from the previous section, we next explore the optimal allocation to the hedge portfolio for individuals with different skill profiles. More importantly, we also assess the value of these hedge portfolios to these individuals.

We assume that individuals have CRRA utility on period consumption, with consumption occurring over T periods with no bequest. We focus on individuals who derive a labor income stream from a single skill j. Equipped with this labor income stream, which they cannot adjust, and an initial amount of financial capital, individuals must decide how much to consume and how much to invest in three financial assets or portfolios: a risk-free asset, the market portfolio, and the hedge portfolio. Formally, their objective is:

$$\max_{\{C_t,\varpi_t,\kappa_t\}_{t=0}^{T}} \mathbb{E}_0\left[\sum_{t=0}^{T}\beta^t\frac{\left(C_t^{1-\gamma}\right)}{1-\gamma}\right], \qquad (4)$$

subject to the intertemporal budget constraint

$$W_{t+1} = (W_t - C_t)[R^F + \varpi_t(R_{t+1}^M - R^F)] + \kappa_t(R_{t+1}^H - 1) + Y_{jt+1} \qquad (5)$$

where C_t is consumption at time t, $\beta \in (0,1)$ is the subjective discount factor, $\gamma > 0$ is the coefficient of relative risk aversion, Y_{jt+1} is the real labor income (from skill j), R^F is the real risk-free return, R^M_{t+1} is the market real return, $(R^H_{t+1} - 1) \equiv r^H_{t+1}$ is the (zero-cost) hedge portfolio's real payoff per unit exposure, ϖ_t is the weight of the market portfolio in the individual's portfolio, and κ_t is the real dollar exposure to the hedge portfolio.[14]

Consistent with our analysis in Section III, we assume that the log real labor income, $y_{t+1} = \log Y_{t+1}$, follows an ARIMA(1,0,1) process:

$$y_{jt+1} = d_j + \phi_j y_{jt} + \iota_{jt+1} + \theta_j \iota_{jt} \qquad (6)$$

where d_j, ϕ_j, and θ_j are the skill-specific parameters of the ARIMA process, and ι_{jt} is the time t (skill j) log real labor income innovation.

In this analysis, we adopt the perspective of an individual standing at the end of our sample period. The individual must decide whether to use the hedge portfolio to hedge human capital innovations and how much capital to commit to it. Only individuals whose skills benefit from hedging (as identified in our in-sample analysis) are assumed to hedge. When doing so, they use the out-of-sample estimates of the mean, variance, and covariance of the hedge portfolio's returns with both the skill innovations and the other financial portfolios. This approach better captures the uncertainty surrounding the hedging strategy.

A. Baseline Model

In our baseline model we concentrate on skill wages. Consistent with the empirical findings in Section III regarding the joint distribution of asset returns and labor income risk, we also assume

[14] The hedge portfolio has its own expected return and risk profile and therefore affects the overall portfolio's expected return and risk profile, but because it costs nothing, it does not consume capital. Formally, its return is given by $r^H_{t+1} = z^{H\prime}_t r_{t+1}$, with weights satisfying $\sum_1^N z^{H\prime}_{t,i} = 0$. Hedge portfolios (weights and returns) vary across skills, but for notational simplicity we omit the skill subscript j.

that all asset returns and labor skill wage innovations are log-normally distributed, with asset returns correlated with three-period-ahead labor income innovations. That is:

$$\begin{bmatrix} \log R_{t+1}^{m} \\ \log R_{t+1}^{h} \\ \iota_{jt+1+} \end{bmatrix} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}), \boldsymbol{\mu} = \begin{bmatrix} \mu_m \\ \mu_h \\ \mu_i \end{bmatrix}, \boldsymbol{\Sigma} = \begin{bmatrix} \sigma_m^2 & \sigma_{mh} & \sigma_{mi} \\ \sigma_{mh} & \sigma_h^2 & \sigma_{hi} \\ \sigma_{mi} & \sigma_{hi} & \sigma_i^2 \end{bmatrix}. \qquad (7)$$

To solve this optimization problem, we adopt the deep learning approach proposed by Duarte et al. (2021). Specifically, we represent closed-loop Markovian controls using a state-dependent neural network and train its parameters—weights and biases—by maximizing the expected cumulative utility along simulated trajectories of the dynamic system described in Equations (5)–(7). For a detailed description, see the Internet Appendix. This approach eliminates the need to explicitly derive and solve the Bellman equation and provides a flexible framework for extending and refining the model.

Table 11 reports results for a range of plausible initial allocations between human capital and financial capital, levels of risk aversion, and investment horizons. Five investment environments are considered. Panel A assumes no access to financial markets, so consumption equals wealth in each period. Panel B allows investment only in the risk-free asset. Panel C permits investment in both the market portfolio and the risk-free asset. Panel D further allows investment in a hedge portfolio constructed from the empirical hedge strategy. Finally, Panel E presents a benchmark case of perfect hedging in which the correlation between the hedge portfolio return and the skill-income innovation is set equal to one.

We find that the welfare gains from access to hedge portfolios, measured as increases in certainty-equivalent consumption, are modest. Certainty equivalents increase by only 0.04% (CRRA = 1) to 0.13% (CRRA = 5) for individuals with access to the hedge portfolio. This reflects the relatively low volatility of the estimated log skill-income innovations and individuals' ability

to smooth consumption through the risk-free asset and the market portfolio, as well as the imperfect nature and volatility of the hedge portfolio itself. Even under perfect hedging of skill-income innovations, the additional welfare benefits remain limited. The certainty-equivalent gains achieved correspond to about 90% of the theoretical gains attainable under perfect hedging when CRRA = 1 and 51% when CRRA = 5.

B. Skill Crashes

We also model the possibility of skill crashes, resulting in the loss of 99% of the skill wage. In particular, let $D_{t+1} \sim \text{Bernoulli}(p)$. In the normal state, if no disaster occurs ($D_{t+1} = 0$), the vector of log returns and innovations follows the original joint normal distribution in (7). However, if a disaster occurs, ($D_{t+1} = 1$), the disaster innovation level equals:

$$\iota_{t+4} = x_{t+4} = \ln(0.01) - d - \theta\iota_{t+3} + (1-\phi)y_{t+3}, \qquad (8)$$

which implies that: $Y_{t+1} = 0.01\, Y_t$.

At the same time, the log return vector follow the conditional distribution implied by that realization:

$$\begin{bmatrix} \log R^M_{t+1} \\ \log R^H_{t+1} \end{bmatrix} \mid D_{t+1} = 1 \sim N(\mu_D, \Sigma_D), \qquad (9)$$

where

$$\mu_D = \begin{bmatrix} \mu_m + \dfrac{\sigma_{mi}}{\sigma_i^2}(x_{t+4} - \mu_i) \\ \mu_h + \dfrac{\sigma_{hi}}{\sigma_i^2}(x_{t+4} - \mu_i) \end{bmatrix},$$

and

$$\Sigma_D = \begin{bmatrix} \sigma_m^2 - \dfrac{\sigma_{mi}^2}{\sigma_i^2} & \sigma_{mh} - \dfrac{\sigma_{mi}\sigma_{hi}}{\sigma_i^2} \\ \sigma_{mh} - \dfrac{\sigma_{mi}\sigma_{hi}}{\sigma_i^2} & \sigma_h^2 - \dfrac{\sigma_{hi}^2}{\sigma_i^2} \end{bmatrix}.$$

As shown in Table 12, allowing for skill crashes leads to significantly more aggressive investments in the hedge portfolio. Allocations to this portfolio increase by at least a factor of four across the specifications we consider. At the same time, the welfare gains associated with the hedge portfolio also rise substantially. These gains now range from 3.95% (CRRA = 1) to 7.36% (CRRA = 5), representing an increase of roughly 50 to 100 times relative to the welfare gains observed in the no–skill-crash scenario.

C. Further Extensions

Several extensions of our model may be worth exploring. A first extension would incorporate retraining dynamics following skill crashes more explicitly. In the current model, innovations to skill-specific wages follow the same stationary stochastic process after the crash. However, one could model such events as a crash after which the stochastic process governing log wage innovations switches to a different regime. For instance, following a crash, the income process could transition to that of a skill with little exposure to AI or technology-related risks. Such a transition would capture the wage dynamics of individuals who retrain and move into occupations that are more neutral or complementary to the new technological environment.

A second extension would introduce idiosyncratic labor income risk. Our empirical skill wage estimates reflect aggregate income processes for individuals possessing a given skill and therefore capture the average income stream associated with that skill. In practice, however, individuals are subject not only to aggregate or skill-specific shocks but also to idiosyncratic income shocks arising from job loss, firm-specific conditions, or individual career trajectories. A natural extension would therefore augment the model with idiosyncratic labor income risk, possibly combined with a deterministic life-cycle component of earnings, as in Catherine (2022).[15]

[15] Catherine's (2022) income process follows closely from the standard MaCurdy (1982) and Abowd and Card (1989) labor income model. One could incorporate the deterministic component of wages, a function of age, and the skewness

Finally, the baseline framework considers the portfolio problem of an individual endowed with a single skill. In reality, individuals possess bundles of skills, and such bundles may themselves provide diversification across technological shocks. A natural extension would therefore model human capital as a portfolio of skills held jointly by the individual. This would allow one to evaluate the extent to which within-person skill diversification mitigates exposure to technological risk and whether these diversification benefits exceed those available through equity markets.

IV. CONCLUSION

Human capital is the largest asset in most individuals' portfolios, yet it is also exposed to significant risk. As artificial intelligence and related technologies reshape the demand for skills, understanding whether financial markets can help hedge the resulting risks becomes increasingly important. This paper studies that question by adopting a granular perspective on human capital, modeling it as a portfolio of skills whose remuneration and demand evolve stochastically over time.

Building on the mimicking portfolio framework of Lamont (2001) and Engle et al. (2020), we design hedge portfolios that load on AI risks, and evaluate their ability to offset innovations in skill-level labor market outcomes. We find that equity portfolios with high AI exposure partially hedge labor market shocks for a nontrivial subset of skills, with economically and statistically significant relationships appearing in both in-sample and out-of-sample tests.

The economic value of these hedging opportunities depends on the nature of human capital risk. When skill income innovations are relatively smooth, welfare gains from access to hedge portfolios are modest. However, when there is a meaningful probability of large skill declines,

of the idiosyncratic income shocks, by relying on estimates from Catherine (2022), Guvenen, Ozkan, and Song (2014), and Guvenen et al. (2021).

such as those that may accompany rapid technological disruption, the value of financial hedges increases substantially.

By showing that equity markets are capable of hedging skill-specific labor income risk, the paper bridges the literatures on human capital risk in labor economics and hedging non-traded risks in asset pricing.

References

Abowd, J. M., & Card, D. (1989). On the Covariance Structure of Earnings and Hours Changes. *Econometrica*, 57(2), 411-445.

Acemoglu, D., & Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In O. Ashenfelter & D. Card (Eds.), *Handbook of labor economics* (Vol. 4B, pp. 1043–1171). Elsevier. https://doi.org/10.1016/S0169-7218(11)02410-5

Acemoglu, D., & Restrepo, P. (2020). Robots and jobs: Evidence from US labor markets. *Journal of Political Economy*, 128(6), 2188-2244.

Akerlof, G. A., Dickens, W. T., & Perry, G. L. (1996). The macroeconomics of low inflation. *Brookings Papers on Economic Activity*, 1996(1), 1–76. https://doi.org/10.2307/2534646

Autor, D. H., Dorn, D., & Hanson, G. H. (2015). Untangling trade and technology: Evidence from local labour markets. *The Economic Journal*, 125(584), 621-646.

Autor, D. H., Levy, F., & Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *Quarterly Journal of Economics*, 118(4), 1279–1333. https://doi.org/10.1162/003355303322552801

Babina, T., Fedyk, A., He, A. X., & Hodson, J. (2023). Firm investments in artificial intelligence technologies and changes in workforce composition (Vol. 31325). National Bureau of Economic Research.

Babina, T., Fedyk, A., He, A., & Hodson, J. (2024). Artificial intelligence, firm growth, and product innovation. Journal of Financial Economics, 153(2), 488–514.

Baker, M. (1997). Growth-rate heterogeneity and the covariance structure of life-cycle earnings. *Journal of Labor Economics*, 15(2), 338-375.

Becker, G. S. (1964). Human capital: A theoretical and empirical analysis with special reference to education. NBER.

Benzoni, L., & Chyruk, O. (2015). The value and risk of human capital. Annual Review of Financial Economics, 7(1), 179-200.

Bewley, T. F. (1999). Why wages don't fall during a recession. Cambridge, MA: Harvard University Press.

Chen, L., Wu, L., & Yang, B. (2025). Surviving the fintech disruption. *Journal of Financial Economics*, 156(2), 391–417.

Cocco, J. F., Gomes, F. J., & Maenhout, P. J. (2005). Consumption and portfolio choice over the life cycle. *The Review of Financial Studies*, 18(2), 491-533.

Davis, S. J., & Willen, P. (2000). Using financial assets to hedge labor income risk: Estimating the benefits. Manuscript, University of Chicago.

Davis, S. J., & Willen, P. (2014). Occupation-level income shocks and asset returns: Their covariance and implications for portfolio choice. *The Quarterly Journal of Finance*, 3(03n04), 1350011.

Deming, D. J. (2017). The growing importance of social skills in the labor market. *Quarterly Journal of Economics*, 132(4), 1593–1640. https://doi.org/10.1093/qje/qjx022

Duarte, V., Fonseca, J., Goodman, A. S., & Parker, J. A. (2021). Simple allocation rules and optimal portfolio choice over the lifecycle (No. w29559). National Bureau of Economic Research.

Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2023). Gpts are gpts: An early look at the labor market impact potential of large language models. arXiv preprint arXiv:2303.10130, 10.

Engle, R. F., Giglio, S., Kelly, B., Lee, H., & Stroebel, J. (2020). Hedging climate change news. *The Review of Financial Studies*, 33(3), 1184-1216.

Frey, C. B., & Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, 114, 254-280.

Grigsby, J., Hurst, E., & Yildirmaz, A. (2021). Aggregate nominal wage adjustments: New evidence from administrative payroll data. *American Economic Review*, 111(1), 428–471. https://doi.org/10.1257/aer.20191178

Guvenen, F. (2009). An empirical investigation of labor income processes. *Review of Economic Dynamics*, 12(1), 58-79.

Guvenen, F., Karahan, F., Ozkan, S., & Song, J. (2021). What do data on millions of US workers reveal about lifecycle earnings dynamics? *Econometrica*, 89(5), 2303-2339.

Guvenen, F., Ozkan, S., & Song, J. (2014). The nature of countercyclical income risk. *Journal of Political Economy*, 122(3), 621-660.

Heaton, J., & Lucas, D. (2000). Portfolio choice and asset prices: The importance of entrepreneurial risk. *The Journal of Finance*, 55(3), 1163-1198.

Hirshleifer, D., Hsu, P. H., & Li, D. (2013). Innovative efficiency and stock returns. *Journal of Financial Economics*, 107(3), 632-654.

Jorgenson, D. W., & Fraumeni, B. M. (1989). The Accumulation of Human and Nonhuman Capital, 1984-84. The Measurement of Saving, Investment, and Wealth, eds. RE Lipsey and HS Tice. Studies in Income and Wealth, 52.

Kahn, S. (1997). Evidence of nominal wage stickiness from microdata. *American Economic Review*, 87(5), 993–1008.

Kochhar, R. (2023). Which US workers are more exposed to AI on their jobs. Pew Research Center, 26.

Kogan, L., Papanikolaou, D., Schmidt, L. D., & Seegmiller, B. (2023). Technology and labor displacement: Evidence from linking patents with worker-level data (No. w31846). National Bureau of Economic Research.

Kogan, L., Papanikolaou, D., Schmidt, L. D., & Song, J. (2021). Technological innovation and labor income risk (No. w26964). National Bureau of Economic Research.

Kogan, L., Papanikolaou, D., Seru, A., & Stoffman, N. (2017). Technological innovation, resource allocation, and growth. *The Quarterly Journal of Economics*, 132(2), 665-712.

Kogan, L., Papanikolaou, D., & Stoffman, N. (2020). Left behind: Creative destruction, inequality, and the stock market. *Journal of Political Economy*, 128(3), 855-906.

Kopeliovich, Y., & Pokojovy, M. (2024). Portfolio optimization with feedback strategies based on artificial neural networks. *Finance Research Letters*, 69, 106185.

Hampole, M., Papanikolaou, D., Schmidt, L. D., & Seegmiller, B. (2025). Artificial intelligence and the labor market (No. w33509). National Bureau of Economic Research.

Lamont, O. A. (2001). Economic tracking portfolios. *Journal of Econometrics*, 105(1), 161-184.

MaCurdy, T. E. (1982). The use of time series processes to model the error structure of earnings in a longitudinal data analysis. *Journal of Econometrics*, 18(1), 83-114.

Mäkelä, E., & Stephany, F. (2024). Complement or substitute? How AI increases the demand for human skills. arXiv preprint arXiv:2412.19754.

Pairolero, N. A., Giczy, A. V., Torres, G., Islam Erana, T., Finlayson, M. A., & Toole, A. A. (2025). The artificial intelligence patent dataset (AIPD) 2023 update. *The Journal of Technology Transfer*, 1-24.

Palacios, M. (2015). Human capital as an asset class: Implications from a general equilibrium model. *The Review of Financial Studies*, 28(4), 978-1023.

Rosen, S. (1974). Hedonic prices and implicit markets: product differentiation in pure competition. *Journal of Political Economy*, 82(1), 34-55.

Stiglitz, J. E. (1984). Theories of wage rigidity (No. w1442). National Bureau of Economic Research.

Tinbergen, J. (1956). On the theory of income distribution. Weltwirtschaftliches archiv, 155-175.

Viceira, L. M. (2001). Optimal portfolio choice for long-horizon investors with nontradable labor income. *The Journal of Finance*, 56(2), 433-470.

Figure 1. Monthly Count of AI News Articles in the Wall Street Journal

This figure plots the monthly number of Wall Street Journal (WSJ) articles that mention at least one of the following keywords: "artificial intelligence," "machine learning," or "deep learning." The sample includes only printed edition articles retrieved from the Factiva database. The vertical red dashed line marks the start of our earliest estimation period (July 2017), and the vertical blue dashed line indicates the beginning of the main sample period (July 2018).

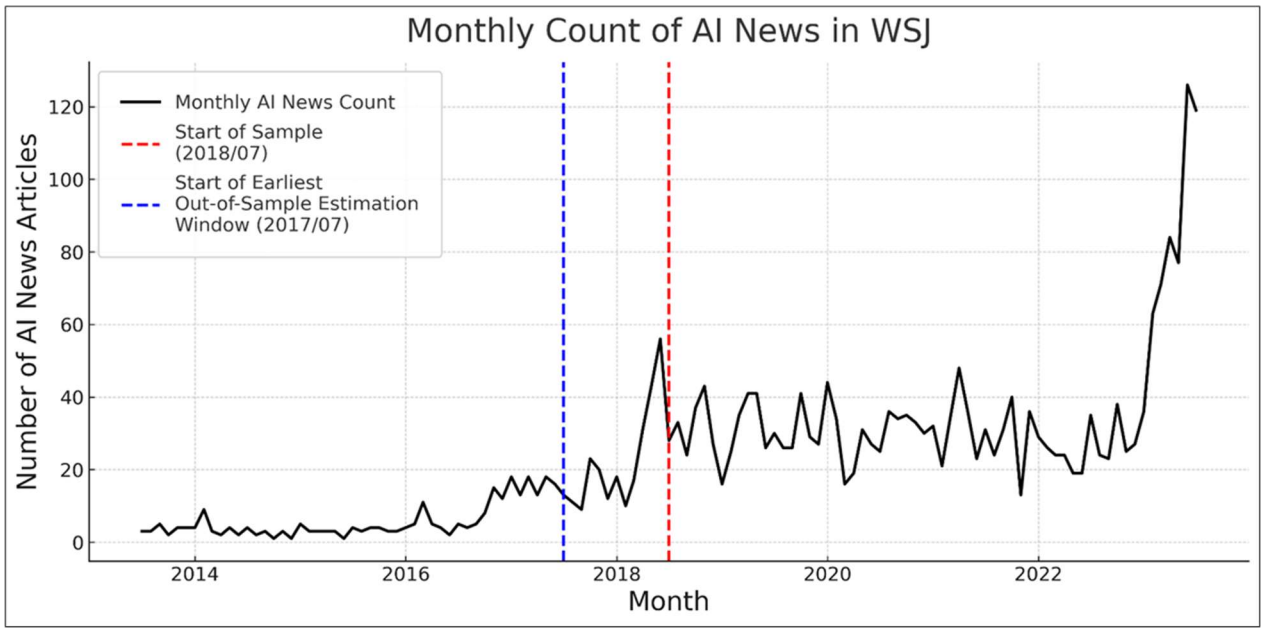


Table 1. Skill Sample Description: Highest and Lowest Average Shadow Wage Skills

This table presents the ten skills with the highest and lowest average shadow wages. For each skill, it reports the skill name, total number of appearances in the sample, its share of total skill appearances across all job postings, the corresponding skill classification, and the average shadow wage over the sample period. Shadow wages are expressed in constant 2018 purchasing-power-equivalent dollars and derived from our log wage model. Skill classifications fall into five categories: Analytical, Social, Managerial, Mechanical, and Fundamental.

10 Skills with the Highest Average Shadow Wage

Skill	Count	% of Sample	Skill Type	Avg Annual Shadow Wage
Registered nurse (rn)	3,459,170	0.95%	Social skills	67,711
CDL class a license	1,536,063	0.42%	Mechanical skills	52,766
Truck driving	1,428,579	0.39%	Mechanical skills	51,884
Negotiation	1,366,992	0.37%	Social skills	45,460
Project management	3,653,640	1.00%	Managerial skills	45,372
Sales	11,513,370	3.16%	Social skills	44,373
Leadership	7,545,639	2.07%	Managerial skills	44,341
Sales prospecting	1,552,767	0.43%	Social skills	43,335
Finance	2,826,476	0.77%	Managerial skills	42,913
Mentorship	1,539,826	0.42%	Social skills	42,819

10 Skills with the Lowest Average Shadow Wage

Skill	Count	% of Sample	Skill Type	Avg Annual Shadow Wage
Merchandising	2,663,989	0.73%	Managerial skills	34,532
Warehousing	4,568,421	1.25%	Mechanical skills	35,298
Retail sales	1,431,329	0.39%	Social skills	35,780
Restaurant operation	2,524,034	0.69%	Managerial skills	35,831
Housekeeping	2,101,920	0.58%	Social skills	36,295
Customer service	19,750,130	5.41%	Social skills	36,567
Cleanliness	2,415,380	0.66%	Mechanical skills	36,822
Lifting ability	6,534,393	1.79%	Mechanical skills	36,842
Data entry	3,068,884	0.84%	Fundamental skills	36,871
Packaging and labeling	1,877,855	0.51%	Mechanical skills	37,375

Table 2. Time Series Properties of Skill Wage, Skill Log Wage, and Skill Count Series

This table summarizes the time-series properties of five skill-level series: the skill wage series, the skill log wage series, the skill count series, and the skill value and log value series, evaluated across the 100 most common skills. For each set of series, the table reports the percentage of skills exhibiting a unit root based on the Augmented Dickey–Fuller (ADF) test at the 5% significance level, the modal ARIMA or SARIMA model selected across skills, and the percentage of skills for which the modal model is the best model.

Skill Wage	
% of skills with unit root (raw series)	31%
Modal ARIMA model	(1, 0, 1)
Modal ARIMA model Percentage	41%
Skill Log Wage	
% of skills with unit root (raw series)	70%
Modal ARIMA model	(1, 0, 1)
Modal ARIMA model Percentage	70%
Skill Count	
% of skills with unit root (raw series)	100%
Modal SARIMA model	(1, 1, 0)x(0, 1, 1, 12)
Modal SARIMA model Percentage	60%
Skill Count × Wage	
% of skills with unit root (raw series)	100%
Modal SARIMA model	(1, 1, 0)x(0, 1, 1, 12)
Modal SARIMA model Percentage	66%
Skill Log Count + Log Wage	
% of skills with unit root (raw series)	99%
Modal SARIMA model	(1, 1, 0)x(0, 1, 1, 12)
Modal SARIMA model Percentage	66%

Table 3. Highest and Lowest Average Wage Occupations

This table presents the ten occupations with the highest and lowest average wages. For each occupation, it reports the occupation name, total number of appearances in the sample, its share of total appearances across all job postings, the corresponding occupation classification, and the average wage over the sample period. Wages are expressed in constant 2018 purchasing-power-equivalent dollars and derived from our log wage model. Occupations are classified into five categories: Analytical, Social, Managerial, Mechanical, and Fundamental.

10 Occupations with the Highest Average Wage

Occupation	Count	% of Sample	Occupation Type	Avg Annual Wage
Nurse Practitioners	230,443	0.43%	Social	102,850
Computer Systems Engineers/Arch.	264,945	0.50%	Analytical	97,241
Software Developers	1,310,419	2.45%	Analytical	95,441
Treasurers and Controllers	184,893	0.35%	Managerial	92,660
Critical Care Nurses	324,419	0.61%	Social	91,209
IT Project Managers	264,271	0.49%	Managerial	88,315
Information Security Analysts	178,471	0.33%	Analytical	86,603
Web Developers	330,704	0.62%	Analytical	85,206
Financial Managers	350,408	0.66%	Managerial	83,823
Database Administrators	178,647	0.33%	Analytical	79,924

10 Occupations with the Lowest Average Wage

Occupation	Count	% of Sample	Occupation Type	Avg Annual Wage
Cashiers	346,907	0.65%	Social	26,760
Dishwashers	210,694	0.39%	Mechanical	26,788
Fast Food and Counter Workers	840,985	1.57%	Social	27,144
Hosts, Restaurants and Coffee Shop	181,942	0.34%	Social	27,508
Waiters and Waitresses	451,368	0.84%	Social	27,988
Janitors and Cleaners	943,651	1.76%	Mechanical	28,150
Hotel, Motel, and Resort Desk Clerks	376,537	0.70%	Social	28,155
Packers and Packagers	209,250	0.39%	Mechanical	28,206
Maids and Housekeeping Cleaners	514,981	0.96%	Social	28,237
Merch. Disp. and Window Trimmers	505,487	0.95%	Social	28,447

Table 4. Time Series Properties of Occupation Wage, Occupation Log Wage, and Occupation Count Series

This table summarizes the time-series properties of five occupation-level series: the occupation wage series, the occupation log wage series, the occupation count series, and the occupation value and log value series, evaluated across the 100 most common occupations. For each set of series, the table reports the percentage of occupations exhibiting a unit root based on the Augmented Dickey–Fuller (ADF) test at the 5% significance level, the modal ARIMA or SARIMA model selected across occupations, and the percentage of occupations for which the modal model is the best model.

Occupation Wage	
% of occupations with unit root (raw series)	6%
Modal ARIMA model	(1, 0, 1)
Modal ARIMA model Percentage	38%
Occupation Log Wage	
% of occupations with unit root (raw series)	38%
Modal ARIMA model	(1, 0, 1)
Modal ARIMA model Percentage	36%
Occupation Count	
% of occupations with unit root (raw series)	99%
Modal SARIMA model	(1, 1, 0)x(0, 1, 1, 12)
Modal SARIMA model Percentage	36%
Occupation Count × Wage	
% of occupations with unit root (raw series)	99%
Modal SARIMA model	(1, 1, 0)x(0, 1, 1, 12)
Modal SARIMA model Percentage	32%
Occupation Log (Count × Wage)	
% of occupations with unit root (raw series)	95%
Modal SARIMA model	(1, 1, 0)x(0, 1, 1, 12)
Modal SARIMA model Percentage	43%

Table 5. Skill Shadow Wage, Count, and Value Hedge Regressions

This table presents summary statistics from regressions of skill shadow wages, posting counts, and values on AI and Fama–French hedge portfolios. Each set comprises 100 regressions corresponding to the 100 most frequent skills. All regressions (one per skill and hedge target) cover the period from July 2018 to July 2023 and incorporate a three-month lag between skill innovations (in wages, counts, or values) and the AI and three EGKLS-style Fama–French portfolios. For each hedge target (skill wages, log wages, counts, values, and log values), the table reports the average R^2 across the 100 regressions, the share of regressions where the AI factor is significant at the 10%, 5%, and 1% levels, and, in square brackets, the corresponding binomial test *p*-values assessing the significance of these proportions. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Innovation	Avg. R^2	AI Portfolio			Three FF Portfolios			All Portfolios		
		Share Sig. (10%)	Share Sig. (5%)	Share Sig. (1%)	Share Sig. (10%)	Share Sig. (5%)	Share Sig. (1%)	Share Sig. (10%)	Share Sig. (5%)	Share Sig. (1%)
Skill Wage	0.117	0.69	0.61	0.29	0.11	0.04	0.00	0.42	0.21	0.05
		[<0.01]***	[<0.01]***	[<0.01]***	[0.42]	[0.74]	[1.00]	[<0.01]***	[<0.01]***	[<0.01]***
Skill Log Wage	0.120	0.73	0.59	0.26	0.23	0.05	0.00	0.43	0.21	0.07
		[<0.01]***	[<0.01]***	[<0.01]***	[<0.01]***	[0.56]	[1.00]	[<0.01]***	[<0.01]***	[<0.01]***
Skill Count	0.106	0.67	0.57	0.31	0.05	0.00	0.00	0.30	0.16	0.04
		[<0.01]***	[<0.01]***	[<0.01]***	[0.98]	[1.00]	[1.00]	[<0.01]***	[<0.01]***	[0.02]**
Skill Count × Wage	0.097	0.59	0.42	0.25	0.05	0.01	0.00	0.26	0.13	0.04
		[<0.01]***	[<0.01]***	[<0.01]***	[0.98]	[0.94]	[1.00]	[<0.01]***	[<0.01]***	[0.02]**
Skill Log Count + Log Wage	0.129	0.70	0.64	0.41	0.16	0.04	0.00	0.49	0.36	0.13
		[<0.01]***	[<0.01]***	[<0.01]***	[0.04]**	[0.74]	[1.00]	[<0.01]***	[<0.01]***	[<0.01]***

Table 6. Occupation Wage, Count, and Value Hedge Regressions

This table presents summary statistics from regressions of occupation wages, posting counts, and values on AI and Fama–French hedge portfolios. Each set comprises 100 regressions corresponding to the 100 most frequent occupations. All regressions (one per occupation and hedge target) cover the period from July 2018 to July 2023 and incorporate a three-month lag between occupation innovations (in wages, counts, or values) and the AI and three EGKLS-style Fama–French portfolios. For each hedge target (occupation wages, log wages, counts, values, and log values), the table reports the average R^2 across the 100 regressions, the share of regressions where the AI factor is significant at the 10%, 5%, and 1% levels, and, in square brackets, the corresponding binomial test *p*-values assessing the significance of these proportions. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Innovation	Avg. R^2	AI Portfolio			Three FF Portfolios			All Portfolios		
		Share Sig. (10%)	Share Sig. (5%)	Share Sig. (1%)	Share Sig. (10%)	Share Sig. (5%)	Share Sig. (1%)	Share Sig. (10%)	Share Sig. (5%)	Share Sig. (1%)
Occupation Wage	0.097	0.32	0.26	0.19	0.19	0.05	0.00	0.31	0.19	0.05
		[<0.01]***	[<0.01]***	[<0.01]***	[<0.01]***	[0.56]	[1.00]	[<0.01]***	[<0.01]***	[<0.01]***
Occupation Log Wage	0.102	0.42	0.36	0.21	0.20	0.07	0.00	0.32	0.19	0.07
		[<0.01]***	[<0.01]***	[<0.01]***	[<0.01]***	[0.23]	[1.00]	[<0.01]***	[<0.01]***	[<0.01]***
Occupation Count	0.101	0.44	0.38	0.21	0.09	0.03	0.00	0.24	0.14	0.05
		[<0.01]***	[<0.01]***	[<0.01]***	[0.68]	[0.88]	[1.00]	[<0.01]***	[<0.01]***	[<0.01]***
Occupation Count × Wage	0.091	0.39	0.28	0.13	0.09	0.05	0.00	0.21	0.13	0.05
		[<0.01]***	[<0.01]***	[<0.01]***	[0.68]	[0.56]	[1.00]	[<0.01]***	[<0.01]***	[<0.01]***
Occupation Log (Count × Wage)	0.132	0.53	0.45	0.33	0.27	0.07	0.00	0.41	0.34	0.17
		[<0.01]***	[<0.01]***	[<0.01]***	[<0.01]***	[0.23]	[1.00]	[<0.01]***	[<0.01]***	[<0.01]***

Table 7. Skills with the Highest and Lowest Exposure to the AI Portfolio

This table lists the skills most frequently ranked within the top 10 or bottom 10 in terms of exposure to the AI portfolio. Exposure is measured using the estimated *t*-statistics from each of our five hedge-target models. For each skill, the table reports its classification (across five categories) and the number of models in which it appears in either the top or bottom 10, regardless of whether the corresponding *t*-statistics are statistically significant, and average *t*-statistics across these models. Only skills that are placed in the top or bottom 10 by at least three of the five models are included.

Skills Most Frequently Ranked in the Top 10 based on AI Portfolio *t*-stats

Skill	Skill Type	# Top 10	Avg. AI *t*-stat
Quality control	Mechanical skills	5	1.562
Packaging and labeling	Mechanical skills	3	0.681
Warehousing	Mechanical skills	3	-0.192
Truck driving	Mechanical skills	3	-0.259
Data analysis	Analytical skills	3	-0.34
Registered nurse (rn)	Social skills	3	-0.343
Computer science	Analytical skills	3	-0.579

Skills Most Frequently Ranked in the Bottom 10 based on AI Portfolio *t*-stats

Skill	Skill Type	# Bottom 10	Avg. AI *t*-stat
Retail sales	Social skills	5	-2.425
Spreadsheets	Analytical skills	4	-2.527
Telephone skills	Social skills	3	-3.958
Multilingualism	Social skills	3	-3.851
Data entry	Fundamental skills	3	-3.789
Enthusiasm	Social skills	3	-1.106

Table 8. Occupations with the Highest and Lowest Exposure to the AI Portfolio

This table lists the occupations most frequently ranked within the top 10 or bottom 10 in terms of exposure to the AI portfolio. Exposure is measured using the estimated *t*-statistics from each of our five hedge-target models. For each occupation, the table reports its classification (across five categories) and the number of models in which it appears in either the top or bottom 10, regardless of whether the corresponding *t*-statistics are statistically significant, and average *t*-statistics across these models. Only occupations that are placed in the top or bottom 10 by at least three of the five models are included.

Occupations Most Frequently Ranked in the Top 10 based on AI Portfolio *t*-stats

Occupation	Occupation Type	# Top 10	Avg. AI *t*-stat
Driver/Sales Workers	Mechanical	3	1.157
Loan Officers	Managerial	3	0.948
Licensed Nurses	Social	3	0.755
Nursing Assistants	Social	3	0.387
Packers and Packagers	Mechanical	3	0.365
Information Security Analysts	Analytical	3	0.291
Critical Care Nurses	Social	3	0.032

Occupations Most Frequently Ranked in the Bottom 10 based on AI Portfolio *t*-stats

Occupation	Occupation Type	# Bottom 10	Avg. AI *t*-stat
Retail Salespersons	Social	5	-2.864
Receptionists and Information Clerks	Social	3	-4.373
Office Clerks, General	Fundamental	3	-4.177

Table 9. Correlation Between Skill Innovations and Hedge Portfolio Out-of-Sample Returns

This table presents the average correlation between skill-level innovations and out-of-sample hedge portfolio returns. Portfolios are constructed using the AI projection portfolio in combination with the three EGKLS-style Fama–French projection portfolios. All projection portfolios are lagged by three months. The hedge portfolios are trained using either a 12-month rolling window (Panel A) or a 36-month rolling window (Panel B). The first part of each panel presents aggregate results for the 100 most frequent skills. The second part presents results for the set of skills that are significant in Table 5. The table shows the skill innovation model (wage, log wage, count, value, and log value), the number of eligible skills for each model, the average out-of-sample correlation between the hedge portfolio returns and skill-level innovations, the share of skills for which the correlation is both positive and statistically significant at the 10%, 5%, and 1% levels, and, in square brackets, the corresponding binomial test p-values assessing the significance of these proportions. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Panel A: All Skills

Innovation	Elegible Skills	Estimation Window	Avg. Correlation	Share Pos. Sig. (10%)	Share Pos. Sig. (5%)	Share Pos. Sig. (1%)
Skill Wage	100	t-13 to t-1	0.178	0.44 [<0.01]***	0.31 [<0.01]***	0.11 [<0.01]***
		t-37 to t-1	0.127	0.26 [<0.01]***	0.12 [<0.01]***	0.02 [0.26]
Skill Log Wage	100	t-13 to t-1	0.155	0.30 [<0.01]***	0.16 [<0.01]***	0.06 [<0.01]***
		t-37 to t-1	0.107	0.15 [0.07]*	0.12 [<0.01]***	0.05 [<0.01]***
Skill Count	100	t-13 to t-1	0.198	0.53 [<0.01]***	0.34 [<0.01]***	0.05 [<0.01]***
		t-37 to t-1	0.116	0.23 [<0.01]***	0.11 [0.01]**	0.01 [0.63]
Skill Count × Wage	100	t-13 to t-1	0.159	0.32 [<0.01]***	0.17 [<0.01]***	0.02 [0.26]
		t-37 to t-1	0.081	0.12 [0.29]	0.02 [0.96]	0.00 [1.00]
Skill Log Count + Log Wage	100	t-13 to t-1	0.052	0.06 [0.94]	0.02 [0.96]	0.01 [0.63]
		t-37 to t-1	0.085	0.14 [0.12]	0.08 [0.12]	0.02 [0.26]

Panel B: Skills for which the Hedge Portfolio is Significant in the In-Sample Analysis

Innovation	Elegible Skills	Estimation Window	Avg. Correlation	Share Pos. Sig. (10%)	Share Pos. Sig. (5%)	Share Pos. Sig. (1%)
Skill Wage	69	t-13 to t-1	0.222	0.55 [<0.01]***	0.41 [<0.01]***	0.16 [<0.01]***
		t-37 to t-1	0.169	0.33 [<0.01]***	0.17 [<0.01]***	0.03 [0.15]
Skill Log Wage	73	t-13 to t-1	0.188	0.37 [<0.01]***	0.22 [<0.01]***	0.08 [<0.01]***
		t-37 to t-1	0.144	0.21 [<0.01]***	0.16 [<0.01]***	0.07 [<0.01]***
Skill Count	67	t-13 to t-1	0.233	0.64 [<0.01]***	0.48 [<0.01]***	0.07 [<0.01]***
		t-37 to t-1	0.162	0.34 [<0.01]***	0.16 [<0.01]***	0.01 [0.49]
Skill Count × Wage	59	t-13 to t-1	0.188	0.42 [<0.01]***	0.24 [<0.01]***	0.02 [0.44]
		t-37 to t-1	0.132	0.20 [0.01]**	0.03 [0.80]	0.00 [1.00]
Skill Log Count + Log Wage	70	t-13 to t-1	0.093	0.09 [0.71]	0.03 [0.87]	0.01 [0.50]
		t-37 to t-1	0.125	0.20 [<0.01]***	0.11 [0.02]**	0.03 [0.15]

Table 10. Correlation Between Occupation Innovations and Hedge Portfolio Out-of-Sample Returns

This table presents the average correlation between occupation-level innovations and out-of-sample hedge portfolio returns. Portfolios are constructed using the AI patent projection portfolio in combination with the three EGKLS-style Fama–French projection portfolios. All factors are lagged by three months. The hedge portfolios are trained using either a 12-month rolling window (Panel A) or a 36-month rolling window (Panel B). The first part of each panel presents aggregate results for the 100 most frequent occupations. The second part presents results for the set of occupations that are significant in Table 6. The table shows the occupation innovation model (wage, log wage, count, value, and log value), the number of eligible occupations for each model, the average out-of-sample correlation between the hedge portfolio returns and occupation-level innovations, the share of jobs for which the correlation is both positive and statistically significant at the 10%, 5%, and 1% levels, and, in square brackets, the corresponding binomial test *p*-values assessing the significance of these proportions. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Panel A: All Occupations

Innovation	Elegible Occupations	Estimation Window	Avg. Correlation	Share Pos. Sig. (10%)	Share Pos. Sig. (5%)	Share Pos. Sig. (1%)
Occupation Wage	100	t-13 to t-1	0.053	0.08	0.05	0.01
				[0.79]	[0.56]	[0.63]
		t-37 to t-1	0.044	0.11	0.08	0.01
				[0.41]	[0.12]	[0.63]
Occupation Log Wage	100	t-13 to t-1	0.049	0.12	0.06	0.01
				[0.29]	[0.38]	[0.63]
		t-37 to t-1	0.038	0.16	0.09	0.03
				[0.04]**	[0.06]*	[0.08]*
Occupation Count	100	t-13 to t-1	0.128	0.30	0.18	0.02
				[<0.01]***	[<0.01]***	[0.26]
		t-37 to t-1	0.103	0.18	0.05	0.01
				[0.01]**	[0.56]	[0.63]
Occupation Count × Wage	100	t-13 to t-1	0.109	0.23	0.15	0.02
				[<0.01]***	[<0.01]***	[0.26]
		t-37 to t-1	0.080	0.12	0.06	0.01
				[0.29]	[0.38]	[0.63]
Occupation Log (Count × Wage)	100	t-13 to t-1	0.088	0.25	0.15	0.05
				[<0.01]***	[<0.01]***	[<0.01]***
		t-37 to t-1	0.100	0.23	0.16	0.05
				[<0.01]***	[<0.01]***	[<0.01]***

Panel B: Occupations for which the Hedge Portfolio is Significant in the In-Sample Analysis

Innovation	Elegible Occupations	Estimation Window	Avg. Correlation	Share Pos. Sig. (10%)	Share Pos. Sig. (5%)	Share Pos. Sig. (1%)
Occupation Wage	32	t-13 to t-1	0.085	0.22	0.16	0.03
				[0.03]**	[0.02]**	[0.27]
		t-37 to t-1	0.111	0.3125	0.25	0.03125
				[<0.01]***	[<0.01]***	0.275
Occupation Log Wage	42	t-13 to t-1	0.070	0.24	0.12	0.02
				[<0.01]***	[0.06]*	[0.34]
		t-37 to t-1	0.097	0.31	0.19	0.07
				[<0.01]***	[<0.01]***	[<0.01]***
Occupation Count	44	t-13 to t-1	0.201	0.52	0.39	0.05
				[<0.01]***	[<0.01]***	[0.07]*
		t-37 to t-1	0.172	0.39	0.11	0.02
				[<0.01]***	0.067*	0.357
Occupation Count × Wage	39	t-13 to t-1	0.186	0.46	0.31	0.05
				[<0.01]***	[<0.01]***	[0.06]*
		t-37 to t-1	0.163	0.31	0.15	0.03
				[<0.01]***	0.012**	0.324
Occupation Log (Count × Wage)	53	t-13 to t-1	0.180	0.47	0.28	0.09
				[<0.01]***	[<0.01]***	[<0.01]***
		t-37 to t-1	0.181	0.43	0.30	0.09
				[<0.01]***	[<0.01]***	[<0.01]***

Table 11. Optimal Consumption and Portfolio Policies Without Skill Crash Risk.

This table reports results for the skill whose innovation has the highest statistically significant correlation with the hedge portfolio in the empirical analysis. The table is based on the baseline simulation environment in which the joint distribution of skill income innovations, market returns, and hedge portfolio returns follows the estimated Gaussian process without rare disaster shocks. The life cycle horizon is T = 240 months. Initial wealth is W0 = 50,000. The annual risk-free rate is set to 1%, and the market portfolio is calibrated to have an annual expected return of 6%. Skill income follows the empirically estimated skill process, and its innovation is jointly simulated with market and hedge portfolio returns using the historical covariance structure described in the simulation procedure. The neural network is trained to maximize expected lifetime CRRA utility for investors with different levels of risk aversion (Gamma = 1, 2, and 5).

Panel A: No financial market access, Ct= Wt

Gamma	EU mean	EU SD	CE mean	CE SD	Cons share mean	Cons share SD	Mkt wt mean	Mkt wt SD	Hedge wt mean	Hedge wt SD
1	1773.87	2.46	3455.23	38.99	1.00	0.00	0.00	0.00	0.00	0.00
2	-0.06	0.00	3424.52	38.85	1.00	0.00	0.00	0.00	0.00	0.00
5	0.00	0.00	3404.70	38.83	1.00	0.00	0.00	0.00	0.00	0.00

Panel B: Risk-free assett Only

Gamma	EU mean	EU SD	CE mean	CE SD	Cons share mean	Cons share SD	Mkt wt mean	Mkt wt SD	Hedge wt mean	Hedge wt SD
1	1784.55	2.31	3628.89	38.50	0.18	0.16	0.00	0.00	0.00	0.00
2	-0.06	0.00	3627.94	38.47	0.18	0.16	0.00	0.00	0.00	0.00
5	0.00	0.00	3625.21	38.40	0.18	0.15	0.00	0.00	0.00	0.00

Panel C: Access to the market portfolio and the risk-free asset

Gamma	EU mean	EU SD	CE mean	CE SD	Cons share mean	Cons share SD	Mkt wt mean	Mkt wt SD	Hedge wt mean	Hedge wt SD
1	1806.21	26.06	4038.04	512.79	0.05	0.08	1.00	0.00	0.00	0.00
2	-0.06	0.00	3898.96	315.20	0.08	0.09	1.00	0.00	0.00	0.00
5	0.00	0.00	3806.00	200.67	0.12	0.12	1.00	0.00	0.00	0.00

Panel D: Access to the market portfolio, the risk-free asset, and the hedge portfolio

Gamma	EU mean	EU SD	CE mean	CE SD	Cons share mean	Cons share SD	Mkt wt mean	Mkt wt SD	Hedge wt mean	Hedge wt SD
1	1806.33	25.86	4039.79	508.97	0.05	0.07	1.00	0.00	-0.26	0.04
2	-0.06	0.00	3900.58	311.35	0.08	0.09	1.00	0.00	-0.27	0.04
5	0.00	0.00	3811.10	198.56	0.12	0.12	1.00	0.00	-0.37	0.14

Panel E: Access to the market portfolio, the risk-free asset, and a perfect hedge portfolio

Gamma	EU mean	EU SD	CE mean	CE SD	Cons share mean	Cons share SD	Mkt wt mean	Mkt wt SD	Hedge wt mean	Hedge wt SD
1	1806.30	26.16	4039.98	515.28	0.05	0.08	1.00	0.00	-0.39	0.04
2	-0.06	0.00	3901.00	305.49	0.08	0.09	1.00	0.00	-0.43	0.04
5	0.00	0.00	3815.91	198.08	0.12	0.12	1.00	0.00	-0.41	0.18

Table 12. Optimal consumption and portfolio policies with Skill Crash Risk.

This table reports results for the skill whose innovation has the highest statistically significant correlation with the hedge portfolio in the empirical analysis. The table is based on a simulation environment that extends the baseline model in Table 1 by introducing rare disaster risk in the skill income process. The life cycle horizon is T = 240 months. Initial wealth is W0 = 50,000. The annual risk-free rate is set to 1%, and the market portfolio is calibrated to have an annual expected return of 6%. In each period, with probability 0.001, a disaster shock occurs that generates a large negative innovation to skill income. Outside the disaster state, the joint distribution of skill income innovations, market returns, and hedge portfolio returns follows the same estimated Gaussian structure used in the baseline model. The neural network is trained to maximize expected lifetime CRRA utility for investors with different levels of risk aversion (Gamma = 1, 2, and 5).

Panel A: No financial market access, Ct= Wt

Gamma	EU mean	EU SD	CE mean	CE SD	Cons share mean	Cons share SD	Mkt wt mean	Mkt wt SD	Hedge wt mean	Hedge wt SD
1	1766.97	14.49	3354.21	209.13	1.00	0.00	0.00	0.00	0.00	0.00
2	-0.08	0.17	3100.38	635.06	1.00	0.00	0.00	0.00	0.00	0.00
5	-220.77	220696.98	2706.14	1341.38	1.00	0.00	0.00	0.00	0.00	0.00

Panel B: Risk-free assett Only

Gamma	EU mean	EU SD	CE mean	CE SD	Cons share mean	Cons share SD	Mkt wt mean	Mkt wt SD	Hedge wt mean	Hedge wt SD
1	1780.41	8.64	3563.01	136.28	0.15	0.13	0.00	0.00	0.00	0.00
2	-0.06	0.00	3545.45	158.98	0.13	0.11	0.00	0.00	0.00	0.00
5	0.00	0.00	3436.25	195.86	0.08	0.02	0.00	0.00	0.00	0.00

Panel C: Access to the market portfolio and the risk-free asset

Gamma	EU mean	EU SD	CE mean	CE SD	Cons share mean	Cons share SD	Mkt wt mean	Mkt wt SD	Hedge wt mean	Hedge wt SD
1	1796.22	29.91	3865.02	540.39	0.07	0.08	1.00	0.00	0.00	0.00
2	-0.06	0.01	3739.33	407.97	0.10	0.11	0.90	0.25	0.00	0.00
5	0.00	0.00	3457.62	199.96	0.09	0.03	0.08	0.01	0.00	0.00

Panel D: Access to the market portfolio, the risk-free asset, and the hedge portfolio

Gamma	EU mean	EU SD	CE mean	CE SD	Cons share mean	Cons share SD	Mkt wt mean	Mkt wt SD	Hedge wt mean	Hedge wt SD
1	1802.98	39.35	4017.83	795.54	0.07	0.10	1.00	0.00	-1.59	0.54
2	-0.06	0.01	3851.65	458.57	0.10	0.12	1.00	0.00	-1.73	0.76
5	0.00	0.00	3712.23	261.93	0.13	0.11	1.00	0.00	-1.74	1.13

Appendix

Appendix Table 1. Largest Average Long and Short Positions in the AI Portfolio

This table shows the industries (2-digit SIC code) with the largest average long and short positions in the AI patent projection portfolio. The portfolio is constructed by weighting firms based on their demeaned cumulative AI patent counts, scaled by each firm's market capitalization. Industries are arranged in descending order of the absolute values of the portfolio weights.

Top Positive Portfolio Weights	
Industry	SIC2
Chemicals and Allied Products	28
Electronic, Electrical Equipment, and Components, except Computer Equipment	36
Measuring, Analyzing, and Controlling Instruments; Photo/Med/Optical Goods; Watches/Clocks	38
Transportation Equipment	37
Industrial and Commercial Machinery and Computer Equipment	35

Top Negative Portfolio Weights	
Industry	SIC2
Electric, Gas, and Sanitary Services	49
Business Services	73
Insurance Carriers	63
Security and Commodiy Brokers, Dealers, Exchange and Services	62
Wholesale Trade - Durable goods	50