

Belief Distortions and Unemployment Fluctuations*

Do Lee

April 22, 2025

Abstract

This paper studies unemployment fluctuations when expectations deviate from a rational benchmark. By using survey forecasts, I decompose time-series and cross-sectional variation in job filling rates. Under subjective beliefs, hiring is driven by predictable errors in cash flow expectations, while discount rates play a minor role. In contrast, rational expectations assign a dominant role to discount rates. Survey measures of the user cost of labor are acyclical, indicating more rigidity under subjective beliefs. The findings suggest that belief distortions can amplify unemployment fluctuations: over-optimism about cash flows during expansions leads to future disappointment, which suppresses labor demand during recessions even as perceived labor costs remain persistently high.

JEL Codes: E7, E24, E27, G4, J64

Keywords: Behavioral, Beliefs, Business Cycles, Unemployment

*New York University, 19 West 4th Street, 6th Floor, New York, NY 10012, dql204@nyu.edu. I am grateful to Jaroslav Borovička, Sophia Chen, Sydney Ludvigson, and Virgiliu Midrigan for their valuable comments. All errors remain my own.

1 Introduction

The search and matching model of the labor market offers a compelling framework for understanding unemployment fluctuations by generating involuntary unemployment through the dynamics of job creation and destruction. Yet a long-standing challenge—known as the “unemployment volatility puzzle”—is that standard search models fail to generate fluctuations in unemployment of the observed magnitude (Shimer, 2005).

One proposed solution introduces rational expectations of time-varying discount rates, linking hiring decisions to fluctuations in risk premia inferred from financial markets (Hall, 2017; Kehoe et al., 2022). The rational expectations assumption of this approach implies that firms correctly understand the relative importance of discount rates and cash flows when posting job vacancies. However, recent work in behavioral finance suggests otherwise: deviations from rational expectations lead investors to over-react to cash flow news, amplifying asset price fluctuations (De La O and Myers, 2021; Bordalo et al., 2024a). While belief distortions are well-documented in financial markets, their consequences for labor markets remain underexplored.

This paper aims to bridge that gap by showing that belief distortions about future cash flows and discount rates can significantly affect unemployment dynamics. I develop a search and matching model where firms’ subjective beliefs may deviate from rational expectations.¹ Because hiring is a forward-looking investment under search frictions, distortions in subjective beliefs can meaningfully impact vacancy postings and job filling rates. From the firm’s hiring condition, I derive a variance decomposition of the job filling rate—a key driver of unemployment fluctuations (Shimer, 2012)—into expectations about future cash flows and discount rates. I estimate this decomposition using survey-based measures of subjective expectations and compare them to machine learning-based proxies for rational expectations.

A central challenge in identifying belief distortions is selecting an appropriate benchmark for rational expectations. Traditional approaches rely on structural models or in-sample estimates from historical data, which either impose strong parametric assumptions or use information not available to firms at decision time (Hansen, 2007; Nagel and Xu, 2021). I address these issues by using machine learning forecasts as an empirical benchmark for undistorted expectations (Bianchi et al., 2022). Machine learning can flexibly process a high-dimensional set of real-time data without imposing a particular theory of belief formation. Systematic deviations from the machine’s predictions then reflect how subjective beliefs overweight or underweight economic signals.

Three main findings highlight the role of belief distortions in labor market dynamics. First, at the aggregate level, subjective expectations about cash flows account for most of variation in job filling rates—94.9% at the five-year horizon—while subjective discount rate expectations contribute only 1.6%. This pattern reverses under rational expectations, where discount rates

¹I use the term “firms’ beliefs” as shorthand to refer to the expectations held by decision makers within firms (Coibion et al., 2018).

dominate. The result suggests that investors under subjective beliefs over-estimate future cash flows during expansions when the job filling rate is low, which consistently leads to disappointment in realized cash flows in the future. Conversely, investors under-estimate the magnitude of fluctuations in future discount rates, where realized future returns following expansions are significantly lower than investors expected. These patterns in job filling rates carry over to unemployment rate decompositions, where subjective expectations similarly amplify the role of cash flow beliefs in driving labor market fluctuations.

Second, cross-sectional analysis at the state and firm levels show that belief distortions in subjective cash flow expectations also drive regional variation in job filling rates, in line with the aggregate level results. Using a Bartik shift-share instrument to isolate exogenous shocks to job filling rates, I show that regions with more volatile subjective cash flow expectations experience larger hiring fluctuations, explaining 62.9% of the variation in regional job filling rates at the 5-year horizon.

Third, subjective expectations of real wages and the user cost of labor are significantly more rigid than their rational counterparts. Search models under rational expectations predict that the user cost of labor responds strongly to business cycle conditions (Kudlyak, 2014; Basu and House, 2016). However, survey data reveal that the user cost of labor is acyclical under subjective expectations. This rigidity can amplify unemployment fluctuations by inducing firms to lay off its workers following adverse shocks instead of adjusting their wages.

Together, these results suggest that belief distortions are a key contributor to unemployment fluctuations. Rational models that emphasize time-varying discount rates may understate the role of distorted cash flow beliefs in driving labor market volatility. The cross-sectional evidence points to a mechanism for persistent regional differences in unemployment through regional differences in the strength of belief distortions. Moreover, rigid wage expectations under subjective beliefs may further amplify unemployment fluctuations by reducing firms' willingness to hire or retain workers during recessions.

Policy interventions targeting unemployment may benefit from addressing systematic biases in investor expectations. During periods of heightened uncertainty, traditional monetary policy tools that focus on adjusting the policy rate may have limited traction if firms base hiring decisions primarily on cash flow expectations rather than discount rates. By expanding forward guidance with clearer communication on the full set of expected economic conditions beyond interest rates, central banks could help stabilize hiring and reduce unemployment volatility.

Related Literature This paper contributes to several strands of literature on unemployment fluctuations, labor market behavior, and expectation formation.

A central challenge in macroeconomics is to explain why unemployment is highly volatile relative to productivity (Shimer, 2005; Ljungqvist and Sargent, 2017). Traditional search and matching models struggle to generate sufficient volatility in unemployment unless firms' responses to shocks are amplified through mechanisms such as rational expectations of time-varying dis-

count rates (Hall, 2017; Borovickova and Borovička, 2017; Kilic and Wachter, 2018; Liu, 2021; Belo et al., 2023; Kehoe et al., 2019; Kehoe et al., 2022; Meeuwis et al., 2023; Mitra et al., 2024). These models assume that firms rationally process information about cash flows and discount rates. My approach of using survey forecasts complements these rational models by introducing subjective expectations. Belief distortions—particularly in expected cash flows—better explain variation in job filling rates, offering an alternative resolution to the unemployment volatility puzzle.

A growing literature embeds non-rational expectations in macro models with labor market frictions (Acharya and Wee, 2020; Mueller et al., 2021; Menzio, 2023; Faberman et al., 2022; Bhandari et al., 2024). Notably, Bhandari et al. (2024) show that systematic pessimism in households and firms amplifies unemployment fluctuations. My paper complements their findings by providing direct survey evidence on the content and cyclicalities of firm expectations, showing that firms over-react to news about earnings expectations while under-reacting to discount rates and labor costs. Cross-sectional analysis shows that regions with more distorted firm beliefs experience larger swings in job filling rates, adding a spatial dimension to belief-driven volatility.

The empirical analysis of this paper builds on existing survey-based evidence on the empirical properties of firm expectations. Ben-David et al. (2013) document persistent over-optimism in CFO forecasts. Gennaioli et al. (2016) document that extrapolative CFO expectations of earnings growth predict corporate investment. Ma et al. (2020) link systematic biases in managerial forecasts to distortions in firm investment. Coibion et al. (2018) and Candia et al. (2020) find that managers' inflation expectations adjust slowly and display substantial dispersion. My paper builds on this work by showing how distortions in survey expectations shape labor markets.

The variance decomposition in this paper builds on recent work using survey-based expectations to reassess the drivers of asset prices. De La O and Myers (2021) show that subjective expectations of cash flow growth, rather than discount rates, explain most of the variation in price-dividend and price-earnings ratios—challenging standard decompositions that assume rational expectations (Campbell and Shiller, 1988; Cochrane, 2007). Bordalo et al. (2024a) find that systematic over-reaction in long-term earnings growth expectations—not time-varying discount rates—accounts for a substantial share of aggregate and cross-sectional return predictability. My paper extends these results to the labor market, showing that biased expectations about firm cash flows not only drive asset prices but also shape real-side decisions such as hiring.

This contribution also relates to a broader literature on over-reaction in subjective expectations (Barberis et al., 1998; Chen et al., 2013; Bordalo et al., 2018; Afrouzi et al., 2023; Bianchi et al., 2023). At the same time, models of inattention and bounded rationality emphasize under-reaction to certain types of information (Mankiw and Reis, 2002; Woodford, 2001; Sims, 2003; Gabaix, 2019). Bianchi et al. (2024b) bridge these perspectives by showing that markets may under-react to news even when investors over-react to all perceived shocks because the over-reaction to multiple perceived signals could offset each other in opposing directions. This insight

is consistent with my finding that firms over-react to news about future earnings but under-react to news about discount rates labor costs.

Informed by this literature, I adopt a machine learning approach to measure rational expectations using a dynamic real-time forecasting framework developed in Bianchi et al. (2022) and Bianchi et al. (2024b). It is based on the principle that rational expectations require agents to efficiently use the full set of real-time information available to them. The algorithm uses high-dimensional prediction models estimated on rolling samples of real-time data to produce a benchmark that is free from human cognitive biases and look-ahead bias, while also addressing overfitting and structural change. The method uses tools from machine learning by training LSTM networks with recursive re-estimation and hyperparameter tuning (Gu et al., 2020, Cong et al., 2020, Bybee et al., 2024). The resulting forecasts are fully ex-ante and provide high-dimensional empirical counterparts to rational expectations for evaluating belief distortions.

The cross-sectional decomposition of this paper contributes to the literature on regional labor market disparities (Sahin et al., 2014; Amior and Manning, 2018; Kehoe et al., 2019; Beraja et al., 2019; Chodorow-Reich and Wieland, 2020). Beraja et al. (2019) show that regional discount factor shocks account for much of the geographic variation in employment during the Great Recession. Kehoe et al. (2019) build on this finding by modeling discount factor shocks as tighter household debt constraints that suppress hiring. On the other hand, Sahin et al. (2014) emphasize that occupational, rather than geographic, mismatch drove much of the post-crisis unemployment rise. My paper complements this work by highlighting a behavioral channel: belief distortions in firms' subjective cash flow expectations can amplify regional hiring fluctuations even absent other structural shocks. This mechanism introduces heterogeneity in expectations as a novel driver of regional business cycles.

The results on subjective wage expectations also relate to debates about wage cyclicality and the user cost of labor (Bils, 1985; Solon et al., 1994; Pissarides, 2009; Gertler and Trigari, 2009; Gertler et al., 2020; Hazell and Taska, 2023). A growing literature emphasizes the user cost of labor—the present value of expected wages to a new hire—as the relevant margin for hiring decisions (Kudlyak, 2014; Basu and House, 2016; Bils et al., 2023). Kudlyak (2014) shows that, under rational expectations, the user cost is far more procyclical than spot or average wages, challenging models that rely on wage rigidity to explain unemployment volatility (Shimer, 2005; Hall, 2005; Blanchard and Galí, 2010; Christiano et al., 2016). I show that the subjective user cost of labor is much more rigid, leading firms to expect persistently high labor costs even in recessions. This belief-driven rigidity suppresses hiring and amplifies unemployment fluctuations through an expectations channel distinct from contractual or institutional wage stickiness.

The rest of the paper proceeds as follows. Section 2 presents a search and matching model with belief distortions and derives a decomposition of the job filling rate. Section 3 describes the data used in the empirical analysis. Section 4 compares the predictive performance of machine and survey forecasts. Section 5 presents the estimated variance decomposition. Section 6 presents

cross-sectional evidence across U.S. states and firms, motivated by a regional extension of the baseline model. Section 7 examines the cyclicity of real wages and the rigidity of subjective wage expectations. Section 8 discusses model extensions and robustness checks to the main result. Section 9 concludes.

2 Theoretical Framework

This section develops a search and matching model of the labor market in which firms' expectations about future cash flows and discount rates may be distorted, leading to fluctuations in job filling rates and unemployment. The model builds on the Diamond (1982), Mortensen (1982), and Pissarides (2009) framework but departs from the standard rational expectations assumption, allowing firms' hiring decisions to be influenced by biased subjective beliefs.²

Environment Consider a discrete time economy populated by a representative household and a representative firm that hires workers in a frictional labor market. The firm uses labor as a single input to production. The household's population is normalized to one and has a continuum of members, where a fraction L_t are employed and the rest are unemployed $U_t = 1 - L_t$. The household's intertemporal consumption decision gives rise to a stochastic discount factor M_{t+1} .

Each period, the firm posts job vacancies at a cost $\kappa > 0$ to maximize its cum-dividend value of equity. We adopt a standard end-of-period matching convention (Petrosky-Nadeau et al., 2018). At the beginning of period t , the stock of employment L_t reflects the total number of workers carried over from the previous period, before any separations or new hires occur during period t . During the period, a fraction δ_t of employed workers separate, while the firm posts vacancies V_t to search for unemployed workers U_t .³ Matches are formed at the end of period t according to a matching function $m(U_t, V_t)$, with job filling rate $q_t \equiv m(U_t, V_t)/V_t$ and job finding rate $f_t \equiv m(U_t, V_t)/U_t$. These new hires enter employment at the start of period $t + 1$, so employment L_t evolves according to a law of motion:

$$L_{t+1} = (1 - \delta_t)L_t + q_t V_t \quad (1)$$

Unemployment $U_t = 1 - L_t$ evolves according to:

$$U_{t+1} = \delta_t(1 - U_t) + (1 - q_t \theta_t)U_t \quad (2)$$

where $\theta_t = V_t/U_t$ denotes labor market tightness, defined as the vacancy-to-unemployment ratio.

I assume that the household owns the equity of the firm and the firm pays out all of its earnings E_t as dividends (Petrosky-Nadeau et al., 2018). I assume that the firm's manager has access to complete markets so that the return to hiring equals the stock market return in equilibrium (Cochrane, 1991).

²See Appendix Section B for more details.

³The job separation rate δ_t may be exogenous or endogenous (e.g., Fujita and Ramey, 2012) because the variance decomposition does not require us to take a stand on the model that generates the job separation rate.

Expectations Let $\mathbb{F}_t[\cdot]$ denote expectations conditional on information available at the beginning of period t , computed based on the firm manager's potentially distorted beliefs. These beliefs may depart from rational expectations $\mathbb{E}_t[\cdot]$ by misweighting information, but I restrict to beliefs that preserve the law of iterated expectations.

Hiring Equation Under search frictions, hiring is forward-looking investment. The firm equates the marginal cost of hiring with the expected discounted value of an additional match:

$$\underbrace{\frac{\kappa}{q_t}}_{\text{Cost of hiring}} = \mathbb{F}_t \left[\underbrace{M_{t+1} \left(\pi_{t+1} + (1 - \delta_{t+1}) \frac{\kappa}{q_{t+1}} \right)}_{\text{Expected discounted value of hiring}} \right] \quad (3)$$

where $\pi_t = \partial \Pi_t / \partial L_t$ is the profit from a marginal worker.⁴ Subjective distortions in beliefs can thus shift the perceived returns to hiring through $\mathbb{F}_t[\cdot]$ and affect equilibrium job filling rates, which in turn affects unemployment through its law of motion in equation (2).

Stock Price Under constant returns to scale (CRS), the firm's equity value per worker is proportional to the marginal value of hiring:⁵

$$\frac{\kappa}{q_t} = \frac{P_t}{L_{t+1}} \quad (4)$$

where employment L_{t+1} is determined at the end of date t under our timing convention from equation (1).⁶ Take logarithms, rearrange terms, and expand the price-employment ratio P_t/L_{t+1} :

$$\log q_t = \log \kappa - \log \left(\frac{P_t}{E_t} \right) + \log \left(\frac{E_{t+1}}{E_t} \right) - \log \left(\frac{E_{t+1}}{L_{t+1}} \right) \quad (5)$$

Defining log price-earnings $pe_t = \log(P_t/E_t)$, log earnings growth $\Delta e_{t+1} = \log(E_{t+1}/E_t)$, and log earnings-employment $el_{t+1} = \log(E_{t+1}/L_{t+1})$, we obtain:

$$\log q_t = \log \kappa - pe_t + \Delta e_{t+1} - el_{t+1} \quad (6)$$

Log-linear Approximation To express the price-earnings ratio pe_t in terms of forward-looking variables, start by log-linearizing the price-dividend ratio $pd_t = \log(P_t/D_t)$ around its long-term average \overline{pd} (Campbell and Shiller, 1988):

$$pd_t = c_{pd} + \Delta d_{t+1} - r_{t+1} + \rho pd_{t+1} \quad (7)$$

⁴The hiring equation is the labor market analogue of the optimality condition for physical capital in the q theory of investment (Hayashi, 1982), where the upfront cost of hiring κ/q_t is analogous to Tobin's marginal q and the separation rate δ_{t+1} is analogous to the depreciation rate. See Lettau and Ludvigson (2002) and Kogan and Papanikolaou (2012) for a similar log-linearization applied for the q theory of physical capital investment.

⁵See Appendix Section B for a derivation (Liu et al., 2009; Kogan and Papanikolaou, 2012).

⁶See Hansen et al. (2005) and Kogan and Papanikolaou (2012) for similar conventions applied for the q theory of investment.

where $r_{t+1} = \log((P_{t+1} + D_{t+1})/P_t)$ is the log stock return (with dividends). c_{pd} and $\rho = \exp(\overline{pd})/(1 + \exp(\overline{pd})) = 0.98$ are constants that arise from the log linearization. Rewrite in terms of price-earnings by using the identity $pe_t = pd_t + de_t$, where de_t is the log payout ratio:

$$pe_t = c_{pd} + \Delta e_{t+1} - r_{t+1} + \rho pe_{t+1} + (1 - \rho)de_{t+1} \quad (8)$$

Since $1 - \rho \approx 0$ and the payout ratio de_t is bounded, $(1 - \rho)de_{t+1}$ can be approximated as a constant (De La O et al., 2024): $c_{pe} \approx c_{pd} + (1 - \rho)de_{t+1}$

$$pe_t \approx c_{pe} + \Delta e_{t+1} - r_{t+1} + \rho pe_{t+1} \quad (9)$$

Recursively substitute for the next h periods:

$$pe_t = \sum_{j=1}^h \rho^{j-1} (c_{pe} + \Delta e_{t+j} - r_{t+j}) + \rho^h pe_{t+h} \quad (10)$$

Decomposition of Job Filling Rates Substitute log-linearized price-earnings from (10) into the hiring equation from (6) to obtain a decomposition of the log job filling rate:

$$\log q_t = c_q + \underbrace{\sum_{j=1}^h \rho^{j-1} r_{t+j}}_{r_{t,t+h}} - \underbrace{\left[el_{t+1} + \sum_{j=2}^h \rho^{j-1} \Delta e_{t+j} \right]}_{e_{t,t+h}} - \underbrace{\rho^h pe_{t+h}}_{pe_{t,t+h}} \quad (11)$$

where $c_q \equiv \log \kappa - \frac{c_{pe}(1-\rho^h)}{1-\rho}$ is a constant. The equation decomposes the job filling rate into future discount rates $r_{t,t+h} \equiv \sum_{j=1}^h \rho^{j-1} r_{t+j}$, cash flows $e_{t,t+h} \equiv el_{t+1} + \sum_{j=2}^h \rho^{j-1} \Delta e_{t+j}$, and price-earnings $pe_{t,t+h} \equiv \rho^h pe_{t+h}$. The cash flow component consists of one period ahead log earnings-employment el_{t+1} , which captures news about short-term cash flows, and $j = 2, \dots, h$ period ahead log earnings growth Δe_{t+j} , which captures long-term news about cash flows.⁷ $pe_{t,t+h}$ is a terminal value that captures other long-term influences beyond h periods into the future not already captured in discount rates and cash flows.

Since equation (11) holds both ex-ante and ex-post, it can be evaluated under either subjective or rational expectations. The *subjective decomposition* replaces ex-post realizations of future outcomes with their subjective expectations $\mathbb{F}_t[\cdot]$:

$$\log q_t = c_q + \underbrace{\sum_{j=1}^h \rho^{j-1} \mathbb{F}_t[r_{t+j}]}_{\mathbb{F}_t[r_{t,t+h}]} - \underbrace{\left[\mathbb{F}_t[el_{t+1}] + \sum_{j=2}^h \rho^{j-1} \mathbb{F}_t[\Delta e_{t+j}] \right]}_{\mathbb{F}_t[e_{t,t+h}]} - \underbrace{\rho^h \mathbb{F}_t[pe_{t+h}]}_{\mathbb{F}_t[pe_{t,t+h}]} \quad (12)$$

⁷The earnings-employment ratio can be interpreted as a measure of the marginal product of labor under constant returns to scale (David et al., 2022).

Alternatively, the *rational decomposition* replaces ex-post realizations of future outcomes with their rational expectations $\mathbb{E}_t[\cdot]$:

$$\log q_t = c_q + \underbrace{\sum_{j=1}^h \rho^{j-1} \mathbb{E}_t[r_{t+j}]}_{\mathbb{E}_t[r_{t,t+h}]} - \underbrace{\left[\mathbb{E}_t[el_{t+1}] + \sum_{j=2}^h \rho^{j-1} \mathbb{E}_t[\Delta e_{t+j}] \right]}_{\mathbb{E}_t[e_{t,t+h}]} - \underbrace{\rho^h \mathbb{E}_t[pe_{t+h}]}_{\mathbb{E}_t[pe_{t,t+h}]} \quad (13)$$

Comparing these decompositions can quantify how belief distortions affect the job filling rate.

Estimation The econometrician can estimate the variance decomposition using predictive regressions of each expected outcome on the current job filling rate. For the subjective decomposition, demean each variable in equation (12), multiply both sides by the current log job filling rate $\log q_t$, and take the sample average:

$$Var[\log q_t] = Cov[\mathbb{F}_t[r_{t,t+h}], \log q_t] - Cov[\mathbb{F}_t[e_{t,t+h}], \log q_t] - Cov[\mathbb{F}_t[pe_{t,t+h}], \log q_t] \quad (14)$$

where $Var[\cdot]$ and $Cov[\cdot]$ are sample variances and covariances based on data observed over a historical sample. Finally, divide both sides by $Var[\log q_t]$ to decompose its variance:

$$1 = \underbrace{\frac{Cov[\mathbb{F}_t[r_{t,t+h}], \log q_t]}{Var[\log q_t]}}_{\text{Discount Rate News}} - \underbrace{\frac{Cov[\mathbb{F}_t[e_{t,t+h}], \log q_t]}{Var[\log q_t]}}_{\text{Cash Flow News}} - \underbrace{\frac{Cov[\mathbb{F}_t[pe_{t,t+h}], \log q_t]}{Var[\log q_t]}}_{\text{Future Price-Earnings News}} \quad (15)$$

The left-hand side represents the full variability in job filling rates, hence is equal to one. Each term on the right reflects the share explained by subjective expectations of discount rates, cash flows, or price-earnings ratios. Under stationarity, the econometrician can estimate these shares using the OLS coefficients from regressing $\mathbb{F}_t[r_{t,t+h}]$, $\mathbb{F}_t[e_{t,t+h}]$, and $\mathbb{F}_t[pe_{t,t+h}]$ on the current log job filling rate $\log q_t$, respectively.

Finally, the decomposition under rational expectations can be estimated similarly based on equation (13) by replacing the subjective expectation $\mathbb{F}_t[\cdot]$ with its rational counterpart $\mathbb{E}_t[\cdot]$. This comparison allows us to assess the role of belief distortions in labor market dynamics and determine whether firms systematically mis-perceive economic conditions when making hiring decisions. The goal is to quantify how biases in subjective beliefs about discount rates and cash flows contribute to fluctuations in the job filling rate.

The variance decomposition has the advantage of not requiring the researcher to take a stand on the deep determinants of job filling rates in any given period. Rather, the evolution of discount rates and cash flows summarize the consequences of the combination of these deep determinants. Not needing to link fluctuations in the job filling rate to a primitive shock makes this approach useful for studying business cycle frequency outcomes that span multiple cycles, each driven by its own unique deep determinants.

3 Data

This section describes the data used to estimate the variance decomposition. For each outcome variable, I use survey forecasts to measure subjective expectations $\mathbb{F}_t[\cdot]$ and machine learning forecasts to measure rational expectations $\mathbb{F}_t[\cdot]$. The final estimation sample is quarterly and spans 2005Q1 to 2021Q4.⁸

Employment Employment L_t is measured using annual total employee counts (EMP) for S&P 500 firms from the CRSP/Compustat Merged Annual Industrial Files. I interpolate this data to a quarterly frequency by using quarterly averages of the fitted values from regressing annual S&P 500 employment on the monthly BLS nonfarm payrolls. Subjective expectations of 1 year ahead annual employment growth are drawn from the CFO survey from 2001Q4 to 2021Q4.

Job Filling Rate Vacancies V_t is measured using JOLTS job openings starting 2000:12 and help-wanted index for earlier periods (Barnichon, 2010). Unemployment U_t is sourced from the BLS unemployment series (UNEMPLOY). The job filling rate q_t is defined as the share of filled vacancies out of unemployment:

$$q_t = \frac{f_t V_t}{U_t} \quad (16)$$

The job finding rate f_t is the share of unemployed workers that find jobs within the period:

$$f_t = 1 - \frac{U_t - U_t^s}{U_{t-1}} \quad (17)$$

where U_t^s is short-term unemployment less than 5 weeks (UEMPLT5). I construct the variables at a monthly frequency, time-aggregate to quarterly averages, and detrend using an HP filter with a smoothing parameter of 10^5 to ensure stationarity (Shimer, 2005).

Stock Returns Stock returns are measured using monthly Center for Research in Security Prices (CRSP) value-weighted returns with dividends (VWRETD). Annualized cumulative h -year log stock returns are compounded from monthly returns: $r_{t,t+h} = (1/h) \sum_{j=1}^{12h} \log(1 + VWRETD_{t+j/12})$. I use the quarterly CFO survey from 2001Q4 to 2021Q4 to construct 1 and 10 year ahead survey forecasts of annualized cumulative log S&P 500 returns. For intermediate horizons between 1 and 10 years, I interpolate linearly between the 1 and 10 year ahead forecasts.

Earnings Quarterly earnings for the S&P 500 are sourced from S&P Global and extended back in time with Shiller’s historical data, yielding a consistent series through 2021Q4.

To construct subjective expectations of future cash flows, I follow the approach of De La O and Myers (2021) and Bordalo et al. (2019), using analyst forecasts from the I/B/E/S database.

⁸See Appendix Section D for details about data sources, and Appendix Section E for details about the machine learning forecasts.

The forecasts reflect the views of financial analysts who actively track firms for their investment research. Analysts have a strong incentive to report their forecasts accurately because these forecasts are not anonymous (Cooper et al., 2001; De La O et al., 2024). Prior research shows that these forecasts are widely followed by market participants and are priced in by investors, lending credibility to their use as proxies for subjective expectations (Kothari et al., 2016).

I/B/E/S provides monthly median analyst forecasts for earnings per share (EPS) at one through four year horizons, as well as long-term growth (LTG) forecasts. The LTG forecast represents the median expected average annual growth rate in operating earnings over the next three to five years. One- through four-year-ahead forward annual log earnings growth forecasts $\mathbb{F}_t[\Delta e_{t+h}]$ for $h = 1, 2, 3, 4$ are log differences between adjacent forecast horizons. For the five-year horizon $\mathbb{F}_t[\Delta e_{t+5}]$, I interpret the LTG forecast as the expected log growth in earnings from year four to five (Bianchi et al., 2024b). The sample spans 1982 to 2021 at a monthly frequency, and the forecasts cover approximately 80% of total market capitalization, providing broad coverage of U.S. public firms.

Price-Earnings Ratio The current price-earnings ratio $PE_t \equiv P_t/E_t$ is calculated using the end-of-quarter S&P 500 stock price index P_t and total earnings E_t . Following De La O and Myers (2021), subjective expectations of log price-earnings $\mathbb{F}_t[pe_{t+h}]$ are derived from the log-linear approximation in equation (10):

$$\mathbb{F}_t[pe_{t+h}] = \frac{1}{\rho^h} pe_t - \frac{1}{\rho^h} \sum_{j=1}^h \rho^{j-1} (c_{pe} + \underbrace{\mathbb{F}_t[\Delta e_{t+j}]}_{\text{Survey (IBES)}} - \underbrace{\mathbb{F}_t[r_{t+j}]}_{\text{Survey (CFO)}}) \quad (18)$$

where expected returns and earnings growth come from the CFO survey and IBES, respectively.

Earnings-Employment Ratio The current earnings-employment ratio is defined as $EL_t \equiv E_t/L_{t+1}$, where E_t denotes quarterly total earnings for the S&P 500 and L_{t+1} is the employment stock at the beginning of period $t + 1$. I measure L_{t+1} using end-of-period employment levels within each quarter. This timing assumption ensures that the measures are consistent with the timing conventions from Section B while still remaining known to firms by the end of period t . Let $el_t = \log(EL_t)$ denote the log earnings-employment ratio. Subjective expectations of future values $\mathbb{F}_t[el_{t+1}] \equiv \mathbb{F}_t[\log E_{t+1} - \log L_{t+1}]$ are constructed as:

$$\mathbb{F}_t[el_{t+1}] = el_t + \underbrace{\mathbb{F}_t[\Delta e_{t,t+1}]}_{\text{Survey (IBES)}} \quad (19)$$

where survey expectations of earnings growth $\mathbb{F}_t[\Delta e_{t,t+h}]$ come from IBES.

Machine Learning Forecasts For each survey forecast, I construct the corresponding machine learning forecast using a Long Short-Term Memory (LSTM) neural network:

$$\mathbb{E}_t[y_{t,t+h}] = G(\mathcal{X}_t, \beta_{h,t}) \quad (20)$$

where $y_{t,t+h}$ denotes the outcome variable y to be predicted h periods ahead of time t . \mathcal{X}_t is a large input dataset of macroeconomic, financial, and textual predictors. $G(\mathcal{X}_t, \beta_{h,t})$ denotes predicted values from Long Short-Term Memory (LSTM) neural networks that can be represented by a potentially high dimensional set of finite-valued parameters $\beta_{h,t}$. The parameters are estimated using an algorithm that takes into account the data-rich environment in which managers operate in (Bianchi et al., 2022 and Bianchi et al., 2024b).

Machine learning forecasts of log price-earnings $\mathbb{E}_t[pe_{t+1}]$ and log earnings-employment $\mathbb{E}_t[el_{t+1}]$ are constructed similarly to the survey counterpart by replacing the survey forecasts on the right-hand side of equation (18) and (A.140) with the corresponding machine learning forecasts.

4 Stylized Facts

Forecasting Performance To assess whether survey respondents systematically misweight relevant information, Table 1 evaluates the out-of-sample accuracy of survey forecasts relative to machine learning forecasts. I measure the relative predictive performance using the ratio of mean-square-forecast-error (MSE) of the machine ($MSE_{\mathbb{E}}$) over that of the survey ($MSE_{\mathbb{F}}$). The out-of-sample R-squared ($OOS R^2$) is the inverse of the MSE ratio, where a positive value indicates better performance of the machine: $1 - MSE_{\mathbb{E}}/MSE_{\mathbb{F}}$. To minimize the influence of random errors in each forecast, the forecast errors are averaged over a sufficiently long out-of-sample testing period spanning from 2005Q1 to 2021Q4. The variables I consider are discount rates $r_{t,t+h}$, cash flows $e_{t,t+h}$, and future price-earnings ratios $pe_{t,t+h}$, as defined in equation (11). These variables influence the job filling rate through the firm’s optimal hiring decision.

Table 1 shows that machine learning forecasts consistently outperform survey forecasts across all variables and horizons. MSE ratios are well below one, and out-of-sample R^2 values are positive, confirming that survey expectations systematically deviate from an unbiased machine learning benchmark. The discrepancy grows with horizon length, indicating that belief distortions in long-term forecasts are larger than those at shorter horizons. These results suggest that subjective expectations used in hiring decisions may be biased in persistent and predictable ways—biases that machine learning models can correct using real-time data. If survey respondents were rational in forming their beliefs, their forecasts would have performed on par with machine learning forecasts. The superior performance of the machine also highlights its ability to process a large amount of real-time data efficiently and objectively, making it a reliable benchmark of undistorted beliefs.

Descriptive Statistics Figure 1 compares subjective and machine expectations for discount rates and cash flows, plotted against the job filling rate. Machine expectations of discount rates exhibit a strong positive relationship with job filling rates, particularly around the Global Financial Crisis. Survey expectations of discount rates, by contrast, are relatively flat and display little sensitivity to the business cycle (Nagel and Xu, 2022). For cash flows, the pattern

Table 1: Accuracy of Machine Learning and Survey Forecasts

Horizon h (Years)	Out-of-sample Testing Period: 2005Q1-2021Q4				
	1	2	3	4	5
(a) Discount Rates $r_{t,t+h}$					
Survey \mathbb{F}	CFO	CFO	CFO	CFO	CFO
$MSE_{\mathbb{E}}/MSE_{\mathbb{F}}$	0.826	0.550	0.352	0.346	0.434
OOS R^2	0.174	0.450	0.648	0.654	0.566
(b) Cash Flows $e_{t,t+h}$					
Survey \mathbb{F}	IBES	IBES	IBES	IBES	IBES
$MSE_{\mathbb{E}}/MSE_{\mathbb{F}}$	0.547	0.304	0.270	0.259	0.235
OOS R^2	0.453	0.696	0.730	0.741	0.765
(c) Price-Earnings $pe_{t,t+h}$					
Survey \mathbb{F}	CFO/IBES	CFO/IBES	CFO/IBES	CFO/IBES	CFO/IBES
$MSE_{\mathbb{E}}/MSE_{\mathbb{F}}$	0.457	0.217	0.188	0.156	0.173
OOS R^2	0.543	0.783	0.812	0.844	0.827

Notes: This table shows the relative mean squared forecast errors (MSE) between machine learning forecasts and survey forecasts. $MSE_{\mathbb{E}}$ and $MSE_{\mathbb{F}}$ denote machine learning and survey mean squared forecast errors, respectively. The out-of-sample R-squared, OOS R^2 , is defined as $1 - MSE_{\mathbb{E}}/MSE_{\mathbb{F}}$, where $MSE_{\mathbb{F}}$ always denotes the MSE of the benchmark forecast. $y_{t,t+h}$ denotes the variable y to be predicted h years ahead of time t . The variables to be predicted are h -year present discounted values of discount rates $r_{t,t+h}$, (negative) cash flows $e_{t,t+h}$, and (negative) price-earnings ratios $pe_{t,t+h}$, as defined in equation (11). Subjective expectations \mathbb{F}_t are based on survey forecasts from the CFO survey for stock returns, and IBES for earnings growth. Machine learning forecasts \mathbb{E}_t are based on a Long Short-Term Memory (LSTM) neural network $G(\mathcal{X}_t, \beta_{h,t})$, whose parameters $\beta_{h,t}$ are estimated in real time using \mathcal{X}_t , a large scale dataset of macroeconomic, financial, and textual data. The out-of-sample testing period is 2005Q1 to 2021Q4.


reverses. Survey expectations show exaggerated cyclical variation, becoming sharply pessimistic during downturns, such as the Global Financial Crisis, when job filling rates are high. Machine forecasts also vary cyclically but to a much lesser extent, indicating that survey respondents tend to over-react to macroeconomic conditions when forming cash flow expectations.

Appendix Table A.1 confirms this visual comparison by summarizing the distributions of survey forecasts and machine learning forecasts for the key components of the variance decomposition. The most striking pattern is the contrast in volatility between survey and machine forecasts. At the five-year horizon, survey-based expectations of discount rates $\mathbb{F}_t[r_{t,t+5}]$ are substantially less volatile than machine-based expectations $\mathbb{E}_t[r_{t,t+5}]$, with standard deviations of 0.073 and 0.129, respectively. In contrast, subjective expectations of cash flows $\mathbb{F}_t[e_{t,t+5}]$ are more volatile than their machine counterparts. At the 5-year horizon, survey expectations of cash flows have a standard deviation of 0.134, compared to 0.198 for machine expectations.

5 Time-Series Decomposition of the Job Filling Rate

The superior forecasting performance of machine learning over survey forecasts suggests the presence of systematic distortions in subjective expectations. This section quantifies how those distortions affect hiring behavior by estimating the contribution of discount rate and cash flow

Figure 1: Job Filling Rates, Discount Rates, and Cash Flows



Notes: Figure plots $h = 5$ year ahead survey forecasts $\mathbb{F}_t[\cdot]$ and machine learning forecasts $\mathbb{E}_t[\cdot]$ of discount rates $r_{t,t+h}$ and cash flows $e_{t,t+h}$ (left axis) against the current detrended job filling rate q_t (right axis). x axis denotes out-of-sample testing period of the forecasts. The forecasts have been aligned with the job filling rate that prevailed at the time of the forecast. Subjective expectations \mathbb{F}_t are based on survey forecasts from the CFO survey for stock returns, and IBES for earnings growth. Machine expectations are based on machine learning forecasts \mathbb{E}_t from Long Short-Term Memory (LSTM) neural networks $G(\mathcal{X}_t, \beta_{h,t})$, whose parameters $\beta_{h,t}$ are estimated in real time using \mathcal{X}_t , a large scale dataset of macroeconomic, financial, and textual data. The out-of-sample forecast testing period is quarterly and spans 2005Q1 to 2021Q4. Cash flow forecasts and job filling rates have been detrend using an HP filter with a smoothing parameter of 10^5 (Shimer, 2005).

expectations to fluctuations in the aggregate job filling rate.

Decomposition Framework Building on the search model from Section 2, the log job filling rate can be written as:

$$\log q_t = c_q + \mathbb{F}_t[r_{t,t+h}] - \mathbb{F}_t[e_{t,t+h}] - \mathbb{F}_t[pe_{t,t+h}] \quad (21)$$

where c_q is a constant and $\mathbb{F}_t[\cdot]$ represents subjective expectations. $r_{t,t+h}$, $e_{t,t+h}$, and $pe_{t,t+h}$ are h -year present discounted values of future discount rates, cash flows, and price-earnings ratios (terminal value), respectively, as defined in equation (11). To quantify the contribution of each term, I estimate variance decompositions by regressing each forecasted component on $\log q_t$:⁹

$$1 = \underbrace{\frac{\text{Cov}[\mathbb{F}_t[r_{t,t+h}], \log q_t]}{\text{Var}[\log q_t]}}_{\text{Discount Rate News}} - \underbrace{\frac{\text{Cov}[\mathbb{F}_t[e_{t,t+h}], \log q_t]}{\text{Var}[\log q_t]}}_{\text{Cash Flow News}} - \underbrace{\frac{\text{Cov}[\mathbb{F}_t[pe_{t,t+h}], \log q_t]}{\text{Var}[\log q_t]}}_{\text{Future Price-Earnings News}} \quad (22)$$

The decomposition under rational expectations $\mathbb{E}_t[\cdot]$ can be estimated similarly while replacing the survey forecast $\mathbb{F}_t[\cdot]$ with its corresponding machine learning counterpart $\mathbb{E}_t[\cdot]$.

Rational Expectations Panel (a) of Table 2 shows that under rational expectations, discount rate news is the dominant driver of variation in job filling rates. At the five-year horizon, rational

⁹The decomposition can also be estimated as a multivariate system by estimating a Vector Autoregression (VAR) of the job filling rate and its key determinants (Table A.2). Since the time series in levels can be serially correlated or nonstationary, I also consider a first differenced decomposition (Table A.3).

Table 2: Time-Series Decomposition of the Job Filling Rate

Horizon h (Years)	1	2	3	4	5
(a) Rational Expectations					
$\log q_t = c_q + \mathbb{E}_t[r_{t,t+h}] - \mathbb{E}_t[e_{t,t+h}] - \mathbb{E}_t[pe_{t,t+h}]$					
Discount Rate	0.187*	0.309**	0.585***	0.653***	0.691***
t -stat	(1.655)	(2.354)	(2.988)	(3.487)	(3.329)
Cash Flow	0.035	0.097	0.123	0.207	0.258
t -stat	(0.097)	(0.361)	(0.461)	(0.785)	(0.918)
Price-Earnings	0.714***	0.513**	0.249	0.199	0.121
t -stat	(2.810)	(2.161)	(0.666)	(0.406)	(0.245)
Residual	0.064	0.081	0.043	-0.059	-0.070
t -stat	(0.143)	(0.212)	(0.086)	(-0.101)	(-0.116)
N	68	68	68	68	68
(b) Subjective Expectations					
$\log q_t = c_q + \mathbb{F}_t[r_{t,t+h}] - \mathbb{F}_t[e_{t,t+h}] - \mathbb{F}_t[pe_{t,t+h}]$					
Discount Rate	0.018	0.018	0.042	0.050	0.016
t -stat	(1.538)	(0.512)	(0.819)	(0.826)	(0.180)
Cash Flow	0.325**	0.641***	0.738***	0.856***	0.964***
t -stat	(1.970)	(4.500)	(4.192)	(4.797)	(7.016)
Price-Earnings	0.629***	0.366	0.207	0.072	0.027
t -stat	(2.794)	(1.410)	(0.765)	(0.187)	(0.077)
Residual	0.027	-0.025	0.014	0.022	-0.007
t -stat	(0.096)	(-0.084)	(0.043)	(0.051)	(-0.017)
N	68	68	68	68	68

Notes: This table reports variance decompositions of the aggregate job filling rate under rational expectations (panel (a)) or subjective expectations (panel (b)). Each row reports the share of the variation in job filling rates that can be explained by h -year expected present discounted values of discount rates $r_{t,t+h}$, (negative) cash flows $e_{t,t+h}$, and (negative) price-earnings ratios $pe_{t,t+h}$, as defined in equation (11). Residual term represents the variation in job filling rates that are not captured by the other components. Subjective expectations \mathbb{F}_t are based on survey forecasts of CFOs and IBES financial analysts. Rational expectations \mathbb{E}_t are based on machine learning forecasts from Long Short-Term Memory (LSTM) neural networks. The sample is quarterly from 2005Q1 to 2021Q4. Newey-West corrected t -statistics with lags = 4 are reported in parentheses: *sig. at 10%. **sig. at 5%. ***sig. at 1%.

discount rates explain 69.1% of the variation in job filling rates, while rational cash flow news accounts for 25.8%.¹⁰ Consistent with the predictions of the search and matching model, higher job filling rates predict either higher future discount rates or lower cash flows. The contribution from terminal price-earnings ratios is small and statistically insignificant at longer horizons. Note that the three components are estimated freely without requiring them to sum to one. Nevertheless, the residual component remains close to zero across horizons, suggesting that the approximations applied to derive the decompositions are reasonably accurate.

These findings align with predictions from rational search-and-matching models that emphasize time-varying risk premia. The large contribution from discount rate news is consistent with rational models that introduce time-varying discount rates to match unemployment fluctuations (Hall, 2017). The increasing importance of discount rate news at longer horizons is consistent

¹⁰First-difference estimates in Table A.3 show broadly similar patterns, with rational cash flow news explaining 58.7% while rational discount rate news explaining only 15.1% of job filling rate variation.

with models that match observed fluctuations in unemployment by modeling hiring as a risky investment with long-duration returns (Kehoe et al., 2022). On the relative importance of risk-free rates and risk premia, Tables A.13 and A.14 show that rational risk-free rate expectations explain little of the variation in the job filling rate. This implies that the explanatory power of discount rate news is driven primarily by the risk premium component, consistent with rational models of labor markets with time-varying risk premia (Borovickova and Borovička, 2017).

Subjective Expectations Panel (b) of Table 2 reveals a striking reversal under subjective expectations. At the five-year horizon, subjective cash flow news explains 96.4% of the variation in job filling rates, while subjective discount rate news accounts for only 1.6%.¹¹ The contribution from subjective price-earnings ratios diminishes with horizon and becomes negligible by year five. These results suggest that firms systematically overweight news about cash flows and underweight news about discount rates. The direction of the relationship is still consistent with the rational case from panel (a), suggesting that mistakes in subjective expectations are about the magnitudes, not about directions.

Compared to the rational benchmark, the implied over-reaction to cash flow news is substantial. Low job filling rates during expansions are associated with a significant disappointment in future cash flows. Defining the survey bias as the difference between subjective and rational expectations, the estimates imply that $96.4\% - 25.8\% = 70.6\%$ of variation in job filling rates can be attributed to the fact that the job filling rate predicts biases in cash flows with a significant positive relationship (Table A.12). These biases capture systematic distortions in survey respondents’ subjective beliefs that the machine learning model could have identified ex-ante.

The residual component remains small under subjective expectations as well, suggesting that any deviations from the model’s approximation are not significantly related to the job filling rate. In other words, survey respondents’ expectations appear internally consistent, but they systematically misperceive which fundamentals matter most for hiring decisions. Any remaining deviations in the residual could be attributed to measurement errors in the survey responses (Ma et al., 2020). The residual term is generally smaller when the decomposition is estimated under subjective expectations, suggesting that incorporating non-rational beliefs into the estimation could improve the ability of the model to explain fluctuations in job filling rates.

Discussion Although the decomposition does not necessarily measure causal relationships, it is useful for quantifying possible sources of variation in the job filling rate. A large estimate for subjective cash flow news means that, whatever shock drives variation in the job filling rate must have a larger impact on subjective cash flow growth expectations than subjective return expectations. Under rational expectations, by contrast, firms correctly interpret those same fluctuations as signals about future risk compensation or discount rates.

¹¹First-difference estimates in Table A.3 show similar results, with subjective cash flow news explaining 83.5% and subjective discount rates explaining only 5.4% of the job filling rate.

This divergence points to a distortion in belief formation: firms over-react to perceived shifts in cash flows while largely ignoring the intertemporal pricing of risk. The result aligns with survey evidence showing that subjective return expectations are acyclical (Nagel and Xu, 2022) or even procyclical (Greenwood and Shleifer, 2014), contrary to the countercyclical discount rate variation implied by rational models (Cochrane, 2017). My findings extend this evidence to labor markets, where hiring decisions appear similarly detached from rational risk-based incentives.

On the other hand, the importance of subjective long-horizon cash flow news in shaping hiring decisions is consistent with models that introduce non-rational expectations about long-term earnings growth to account for fluctuations in asset prices (Bordalo et al., 2024a) and the business cycle (Bordalo et al., 2024b). Specifically, the result is consistent with investor over-reaction to fluctuations in the payout share of output (Bianchi et al., 2024b). My results suggest that firms exhibit similar behavior when hiring: they respond primarily to exaggerated expectations about future earnings, not shifts in discount rates. The increasing importance of subjective cash flow news at longer horizons is consistent with behavioral models in which over-reaction increases with the horizon of the survey forecast (Bordalo et al., 2020; Bianchi et al., 2024a; Augenblick et al., 2024).

While this paper focuses on publicly listed firms, preliminary evidence suggest that similar patterns emerge among smaller private businesses. A 2010 report from the National Federation of Independent Business (NFIB) on small business credit during the recession shows that hiring decisions were primarily driven by pessimism about future sales rather than financing constraints. At the time, 51% of small employers cited weak sales expectations as their top concern, compared to just 8% who cited access to credit (Dennis, 2010). This suggests that subjective beliefs about future cash flows—not discount rates—also shape employment decisions in the small business sector.

These belief distortions offer a new explanation for unemployment volatility. Firms that misinterpret hiring conditions as news about cash flows—rather than changes in risk—may cut vacancies excessively in downturns, amplifying labor market fluctuations beyond what rational models predict. The next section quantifies this mechanism by directly decomposing the unemployment rate.

Implications for Unemployment Fluctuations To complement the decomposition of the job filling rate, this section analyzes the unemployment rate directly. While the job filling rate captures the proximate drivers of employment dynamics in search models, the unemployment rate is the key macroeconomic outcome of interest and the direct target of policy.

A direct decomposition of unemployment offers three advantages. First, it provides a more interpretable measure of the macroeconomic significance of belief distortions. Second, it captures how these distortions propagate through the nonlinear law of motion for unemployment, which also depends on job separations and labor market tightness. Third, it links observed unemployment fluctuations to distortions in the underlying expectations.

Start from the unemployment accumulation equation:

$$U_{t+1} = \delta_t(1 - U_t) + (1 - q_t\theta_t)U_t \quad (23)$$

For some variable X_t , denote the steady state values without time subscripts X , and define log deviations from steady state as $\hat{X}_t = \log(X_t) - \log(X)$. Log-linearize the unemployment accumulation equation around the steady state (See Appendix Section C for a derivation):

$$\hat{U}_{t+1} \approx q\theta\hat{\delta}_t + (1 - \delta - q\theta)\hat{U}_t - q\theta\hat{q}_t - q\theta\hat{\theta}_t \quad (24)$$

Finally, substitute in equation (11), which is a decomposition of the job filling rate \hat{q}_t into discount rate, cash flow, and future price-earnings components:

$$\hat{U}_{t+1} \approx - \underbrace{q\theta \cdot \hat{r}_{t,t+h}}_{\text{Discount Rate}} + \underbrace{q\theta \cdot \hat{e}_{t,t+h}}_{\text{Cash Flow}} + \underbrace{q\theta \cdot \hat{p}e_{t,t+h}}_{\text{Future Price-Earning}} + \underbrace{(1 - \delta - q\theta) \cdot \hat{U}_t - q\theta \cdot \hat{\theta}_t + q\theta \cdot \hat{\delta}_t}_{\text{Lag Unemployment, Tightness, Separations}} \quad (25)$$

This equation can be evaluated using either subjective $\mathbb{F}_t[\cdot]$ or rational $\mathbb{F}_t[\cdot]$ expectations. I estimate it using regressions in first differences to ensure stationarity and identify the relative contributions of each component to observed unemployment fluctuations.

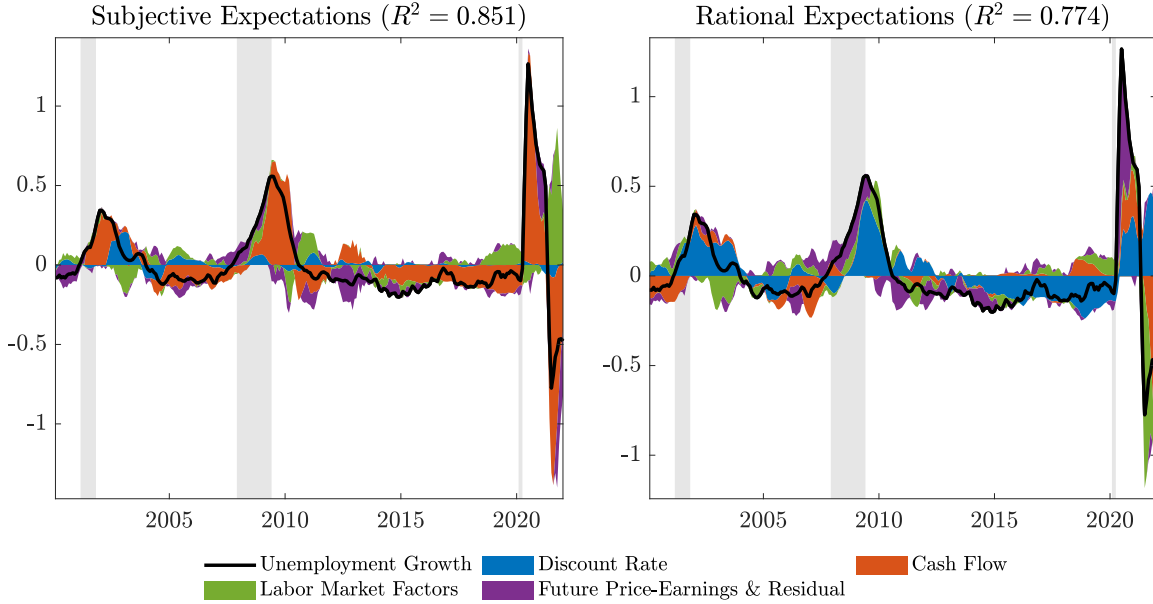
Figure 2 plots the actual annual change in unemployment against its model-implied decomposition using both subjective and rational expectations. Under subjective beliefs, fluctuations in unemployment closely track the component attributed to expected cash flow news. The model captures the sharp rise and fall in unemployment during the COVID-19 crisis with considerable precision. The fit of the model in terms of R^2 is notably better under subjective expectations than under rational expectations, indicating that incorporating belief distortions improves the model’s explanatory power.

This finding underscores that belief distortions can have macroeconomically meaningful consequences. Systematic misperceptions about future cash flows—rather than discount rates or job separations—can significantly influence firms’ hiring behavior and, in turn, unemployment dynamics. The resulting gap between unemployment predicted under rational versus subjective expectations can be interpreted as a reduced-form belief-driven unemployment wedge. This wedge provides a novel empirical counterpart to theoretical mechanisms in existing behavioral models where non-rational beliefs amplify business cycle volatility (Bordalo et al., 2021; Bianchi et al., 2023). By propagating through forward-looking hiring behavior, belief distortions lead not just to forecast errors, but to real fluctuations in labor market outcomes.

6 Cross-Sectional Evidence

The aggregate analysis in Section 5 shows that belief distortions in subjective expectations play an important role in explaining hiring fluctuations. This section extends that analysis by exploiting cross-sectional variation in firm- and region-level data to strengthen identification and test whether the theoretical mechanism generalizes beyond aggregate dynamics.

Figure 2: Time-Series Decomposition of the Unemployment Rate



Notes: Figure plots decompositions of log annual growth in the unemployment rate from equation (25), under subjective (left panel) or rational (right panel) expectations. Labor market factors include lagged unemployment, labor market tightness and job separations. Residual term represents the variation in job filling rates that are not captured by the other components. Subjective expectations \mathbb{E}_t are based on survey forecasts from CFOs and IBES financial analysts. Rational expectations \mathbb{E}_t are based on machine learning forecasts from Long Short-Term Memory (LSTM) neural networks. NBER recessions are shown with gray shaded bars.

Overview While the aggregate-level variance decompositions are informative, they cannot establish causality, and the limited number of business cycles in the time series restricts inference (Chodorow-Reich and Wieland, 2020). This section addresses these challenges by extending the aggregate model to a regional framework. In estimating the regional model, I introduce a Bartik shift-share instrument for the job filling rate to address endogeneity challenges in identifying the relative importance of subjective discount rate and cash flow expectations.

Specifically, I investigate whether regional labor markets characterized by more distorted subjective cash flow expectations experience larger swings in job filling rates. This analysis is motivated by empirical evidence of substantial geographic variation in unemployment dynamics, especially during crises (Beraja et al., 2019, Kehoe et al., 2019; Chodorow-Reich and Wieland, 2020). While existing work studies these regional differences under a rational expectations framework, regional heterogeneity in subjective beliefs may also be an important explanatory factor.

Regional Model To guide the empirical strategy, I extend the baseline search model to a multi-region, multi-sector environment, building from the models in Kehoe et al. (2019) and Chodorow-Reich and Wieland (2020). The economy consists of a continuum of islands indexed by s . Each island produces a differentiated variety of tradable goods that is consumed everywhere and a nontradable good. Both of these goods are produced using intermediate goods. Each consumer is endowed with one of two types of skills which are used in different intensities in

the nontradable and tradable goods sectors. Labor is immobile across islands but can switch sectors.¹² Consumers receive utility from a composite consumption good that is either purchased in the market or produced at home. Consumers and firms are ex-ante homogeneous and share the same subjective expectation $\mathbb{F}_t[\cdot]$. The islands only differ in the shocks that hit them.

In this environment, the hiring decision for a firm producing intermediate good i in island s satisfies a regional analog of the aggregate hiring equation:¹³

$$\frac{\kappa}{q_{s,i,t}} = \mathbb{F}_t \left[M_{s,t,t+1} \left(\pi_{s,i,t+1} + (1 - \delta_{s,i,t+1}) \frac{\kappa}{q_{s,i,t+1}} \right) \right] \quad (26)$$

where $q_{s,i,t}$ is the regional job filling rate, κ is the vacancy posting cost, δ is the job separation rate, $M_{s,t,t+1}$ is the stochastic discount factor, and $\pi_{s,i,t+1}$ is the marginal profit of a new hire.

Decomposition of Regional Job Filling Rates Taking logs and applying the decomposition developed in Section 2, the job filling rate can be written as:

$$\log q_{s,i,t} = c_q + \underbrace{\sum_{j=1}^h \rho^{j-1} r_{s,i,t+j}}_{r_{s,i,t,t+h}} - \underbrace{\left[e l_{s,i,t+1} + \sum_{j=2}^h \rho^{j-1} \Delta e_{s,i,t+j} \right]}_{e_{s,i,t,t+h}} - \underbrace{\rho^h p e_{s,i,t+h}}_{p e_{s,i,t,t+h}} \quad (27)$$

where c_q is a constant. This regional decomposition mirrors the aggregate decomposition from equation (11). The equation implies that if firms form distorted beliefs about future returns or earnings, those distortions should influence hiring decisions at the regional level as well. Since the decomposition holds both ex-ante and ex-post, I consider an ex-ante decomposition under subjective expectations $\mathbb{F}_t[\cdot]$:

$$\log q_{s,i,t} = c_q + \mathbb{F}_t[r_{s,i,t,t+h}] - \mathbb{F}_t[e_{s,i,t,t+h}] - \mathbb{F}_t[p e_{s,i,t,t+h}] \quad (28)$$

where $\mathbb{F}_t[x_{s,i,t,t+h}]$ denotes the h period ahead time t subjective expectations of variable x for a firm producing intermediate good i in island s . Equation (28) implies the following variance decomposition of the regional job filling rate:

$$1 = \underbrace{\frac{Cov[\mathbb{F}_t[r_{s,i,t,t+h}], \log q_{s,i,t}]}{Var[\log q_{s,i,t}]}_{\text{Discount Rate News}}} - \underbrace{\frac{Cov[\mathbb{F}_t[e_{s,i,t,t+h}], \log q_{s,i,t}]}{Var[\log q_{s,i,t}]}_{\text{Cash Flow News}}} - \underbrace{\frac{Cov[\mathbb{F}_t[p e_{s,i,t,t+h}], \log q_{s,i,t}]}{Var[\log q_{s,i,t}]}_{\text{Future Price-Earnings News}}} \quad (29)$$

where $Var[\cdot]$ and $Cov[\cdot]$ are sample variances and covariances based on data observed over a historical sample. The left-hand side represents the full variability in regional job filling rates, hence is equal to one. Each term on the right reflects the share explained by subjective expectations of discount rates, cash flows, or price-earnings ratios. Under stationarity, the econometrician can estimate these shares from regressing $\mathbb{F}_t[r_{s,i,t,t+h}]$, $\mathbb{F}_t[e_{s,i,t,t+h}]$, and $\mathbb{F}_{s,i,t}[p e_{t,t+h}]$ on the current log regional job filling rate $\log q_{s,i,t}$, respectively. Since the regional panel data can be serially correlated or nonstationary in levels, I also consider a first differenced decomposition.

¹²This assumption aligns with empirical evidence indicating that labor markets are predominantly local in nature (Manning and Petrongolo, 2017).

¹³See Section C.1 for a derivation.

Empirical Specification: OLS The goal of the regional decomposition is to assess whether perceived shocks to discount rates or cash flows drive variation in job filling rates at the local labor market level. I begin by estimating the relationship between firm-level subjective expectations and a direct measure of state-level job filling rates. The job filling rate in state s is computed as a vacancy-weighted average across industries:

$$q_{s,t} = \sum_i \left(\frac{V_{s,i,t}}{\sum_{i'} V_{s,i',t}} \right) q_{s,i,t} \quad (30)$$

where s is measured as U.S. states, time t is measured at monthly frequency, $i = 1, \dots, I$ are 2-digit NAICS industries. $V_{s,i,t}$ and $q_{s,i,t}$ denote vacancies and job filling rates in island s and industry i , respectively. I use JOLTS data to construct regional job filling rates $q_{s,t}$ as the ratio of total hires to job openings at the state level (Leduc and Liu, 2020). Using this measure, I estimate the following firm-level regression:

$$\mathbb{F}_t[y_{f,s,t,t+5}] = \beta \cdot q_{s,t} + \alpha_f + \alpha_t + \gamma' X_{f,s,t} + \varepsilon_{f,s,t}, \quad y = r, e, pe \quad (31)$$

for firm f , state s , and time t . Following Korniotis (2008), each firm is mapped to the state in which the firm has its headquarter. The parameter of interest β captures the share of variation in the measure of job filling rate that can be explained by $h = 5$ year ahead subjective expectations about discount rates r , cash flows e , or future price-earnings pe . α_f are firm fixed effects, and α_t are time fixed effects. $X_{f,s,t}$ is a vector of controls including decile dummies for firm size measured by number of employees at the time of the survey forecast and a linear trend.

Empirical Specification: Bartik Shift-Share Instrument A key challenge in estimating the variance decomposition is that subjective expectations and job filling rates may be jointly determined—for example, if firms revise expectations in response to local conditions. I address the challenge by constructing a Bartik shift-share instrument $\tilde{q}_{i,t}$ which provides an exogenous source of shocks to the regional labor market:

$$\tilde{q}_{s,t} = \sum_{i \in I} \omega_{s,i,t} q_{i,t}, \quad \omega_{s,i,t} = \frac{L_{s,i,t}}{\sum_{i'} L_{s,i',t}} \quad (32)$$

where time t is measured at monthly frequency, and $i = 1, \dots, I$ are 2-digit NAICS industries (excluding public administration). $\omega_{s,i,t}$ is the employment share of industry i in state s in the year prior to time t , sourced from the BLS Quarterly Census of Employment and Wages (QCEW). I measure industry-wide job filling rates $q_{i,t}$ as the ratio between total hires and job vacancies at the industry level, both from JOLTS (Leduc and Liu, 2020). The summation over i aggregates these industry-level job filling rates using each state's industry weights. The instrument exploits national industry-level demand shocks that differentially affect regional labor markets based on their industry composition.

Using the Bartik instrument for regional labor market tightness, I estimate the following regression at the firm-level:

$$\mathbb{F}_t[y_{f,s,t,t+5}] = \beta \cdot \tilde{q}_{s,t} + \alpha_f + \alpha_t + \gamma' X_{f,s,t} + \varepsilon_{f,s,t}, \quad y = r, e, pe \quad (33)$$

where $\tilde{q}_{s,t}$ is the Bartik shift-share instrument for regional job filling rates from equation (32), and the other terms are the same as in the OLS specification from (31). The parameter of interest β captures the share of variation in the shift-share instrumented job filling rate that can be explained by $h = 5$ year ahead subjective expectations about discount rates r , cash flows e , or future price-earnings pe .

Identification Assumptions Compared to the direct OLS approach, the Bartik specification provides more credible identification for two reasons. First, state-level hiring and vacancy data may suffer from measurement error. Second, subjective expectations and observed labor market conditions may be endogenously linked. If firms revise beliefs in response to local shocks, the OLS coefficient may reflect a feedback loop that overstates the role of beliefs in driving hiring.

The shift-share instrument helps isolate plausibly exogenous variation by using industry-level trends that are predetermined with respect to firm-level expectations. Since the instrument aggregates national industry trends, it is predetermined with respect to firm-level expectations formed at time t . The use of aggregate industry-level variation can also control for the possibility that firms in each region have private information that affect their hiring decisions. The identifying assumption is that conditional on fixed effects and controls, there are no omitted factors that simultaneously affect both the national industry hiring trends and the subjective expectations of firms located in states more exposed to those industries.

While the shift-share approach strengthens causal interpretation, it is not immune to potential threats. A leading concern is the possibility of differential pre-trends. States with industry mixes that make them more exposed to national shocks may have had systematically different trajectories even before the shock occurred. To mitigate this concern, I include a comprehensive set of controls. Because larger firms tend to pay higher wages, controlling for firm size helps account for differences in human capital and wage premia across firms (Oi and Idson, 1999; Elsby and Michaels, 2013). Firm fixed effects α_f absorb time-invariant characteristics such as productivity (Acemoglu and Hawkins, 2014) and matching rates (Kaas and Kircher, 2015). Time fixed effects α_t account for aggregate shocks, national policy changes, and long-run expectations that are common across regions (Gavazza et al., 2018; Hazell et al., 2022).

Cross-Sectional Decomposition of Regional Job Filling Rates Table 3 reports the estimated decompositions. The raw regression coefficients reported in the table can be converted into a variance decomposition by taking the negative of the coefficients for cash flow and price-earnings expectations, since the negative sign on these coefficients reflect the inverse relationship with job filling rates in the search model.

Table 3: Cross-Sectional Decomposition of the Job Filling Rate

(a) Subjective Expectations: Baseline						
$\mathbb{E}_t[y_{f,s,t,t+5}] = \beta \cdot q_{s,t} + \alpha_f + \varepsilon_{f,s,t}, \quad y = r, e, pe$						
	(1)	(2)	(3)	(4)	(5)	(6)
	Discount Rate	Cash Flow	Price Earning	Discount Rate	Cash Flow	Price Earning
Job Filling Rate	0.1623*** (0.0415)	-0.6290*** (0.0609)	-0.2172*** (0.0631)	0.1002 (0.0988)	-0.7090*** (0.1030)	-0.1938** (0.0703)
Observations	611,636	611,636	611,636	549,513	549,513	549,513
Adj. R^2	0.49	0.90	0.48	0.03	0.09	0.06
Specification	Levels	Levels	Levels	Differences	Differences	Differences
Instrument	OLS	OLS	OLS	OLS	OLS	OLS
Controls	No	No	No	No	No	No
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	No	No	No	No	No	No
(b) Subjective Expectations: Time Fixed Effects						
$\mathbb{E}_t[y_{f,s,t,t+5}] = \beta \cdot q_{s,t} + \alpha_f + \alpha_t + \gamma' X_{f,s,t} + \varepsilon_{f,s,t}, \quad y = r, e, pe$						
	(1)	(2)	(3)	(4)	(5)	(6)
	Discount Rate	Cash Flow	Price Earning	Discount Rate	Cash Flow	Price Earning
Job Filling Rate	0.0603 (0.3296)	-0.5644*** (0.1469)	-0.0527 (0.1997)	0.1164 (0.1124)	-0.5478*** (0.1397)	-0.0269 (0.0553)
Observations	611,636	611,636	611,636	549,513	549,513	549,513
Adj. R^2	0.52	0.91	0.49	0.02	0.12	0.00
Specification	Levels	Levels	Levels	Differences	Differences	Differences
Instrument	OLS	OLS	OLS	OLS	OLS	OLS
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
(c) Subjective Expectations: Bartik Instrument						
$\mathbb{E}_t[y_{f,s,t,t+5}] = \beta \cdot \tilde{q}_{s,t} + \alpha_f + \alpha_t + \gamma' X_{f,s,t} + \varepsilon_{f,s,t}, \quad y = r, e, pe$						
	(1)	(2)	(3)	(4)	(5)	(6)
	Discount Rate	Cash Flow	Price Earning	Discount Rate	Cash Flow	Price Earning
Job Filling Rate	0.0241 (0.2090)	-0.3764*** (0.1307)	-0.0180 (0.0463)	0.0457 (0.0687)	-0.3227*** (0.0939)	-0.0019 (0.0306)
Observations	611,636	611,636	611,636	549,513	549,513	549,513
Adj. R^2	0.52	0.91	0.91	0.02	0.12	0.01
Specification	Levels	Levels	Levels	Differences	Differences	Differences
Instrument	Bartik	Bartik	Bartik	Bartik	Bartik	Bartik
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table estimates the firm-level relationship between subjective expectations of discount rate r , cash flow e , or price-earnings pe against regional job filling rates for the state in which the firm has its headquarter. Columns (1)–(3) present estimates in levels. Columns (4)–(6) present estimates in first differences. Panel (a) reports OLS regressions with no other controls. Panel (b) introduces firm-level controls and time fixed effects. Panel (c) reports results using a Bartik shift-share instrument $\tilde{q}_{s,t}$, which interacts state-level employment shares in 2-digit NAICS industries with national industry-specific job filling rate fluctuations. Subjective expectations \mathbb{E}_t are based on survey forecasts of CFOs and IBES financial analysts. $X_{f,s,t}$ is a vector of controls including decile dummies for firm size and a linear trend. Sample is monthly from 2001M1 to 2021M12. Observations are weighted by each firm's market capitalization. Standard errors clustered by state are reported in parentheses: *sig. at 10%. **sig. at 5%. ***sig. at 1%.

Panel (a) presents baseline OLS results. At the five-year horizon, subjective cash flow expectations account for 62.9% of the variation in regional job filling rates, while subjective discount rates and price-earnings expectations explain 16.2% and 21.7%, respectively. Results using first differences yield similar patterns. Panel (b) includes time fixed effects and additional controls. Although the contribution of subjective cash flow expectations declines slightly to 56.4%, it remains the dominant driver. Panel (c) uses the Bartik shift-share instrument for job filling rates. Even under this stricter identification, subjective cash flow expectations explain 37.6% of the variation in levels and 32.3% in differences. The shift-share estimates are likely smaller than their OLS counterpart because the instrument only captures only the shocks that are exogenous to the regional labor market. The muted estimate from using the instrument suggests that state specific factors such as shifts in regional labor supply could amplify the role of subjective beliefs, since the larger estimates from using the OLS specification capture these alternative channels.

Discussion The results suggest that distorted cash flow expectations are an important driver of regional labor market volatility. The strong regional co-movement between cash flow expectations and job filling rates remains robust when instrumented using a Bartik shift-share, supporting a causal interpretation. Regions where firms over-react more strongly to cash flow news experience deeper hiring cuts during downturns and over-hiring during expansions, amplifying the volatility of regional business cycles. These findings suggest that persistent differences in unemployment across regions may reflect not only structural factors—such as industry composition or labor supply—but also variation in how firms form and act on expectations.

The variance share attributed to cash flow beliefs is somewhat smaller in the cross-section than in the aggregate time series. The large aggregate estimate on subjective cash flow news is consistent with models of heterogeneous firms where the dispersion in the marginal product of capital or labor amplify fluctuations in aggregate total factor productivity (Ma et al., 2020; David et al., 2022; Ropele et al., 2024). When MPL dispersion increases, resources are misallocated, lowering aggregate output per unit of labor. Such dispersion can influence aggregate output because search frictions in the labor market or wage rigidity prevents workers from adjusting immediately to the shock by moving across regions or industries. Consistent with the model’s predictions, Figure A.1 shows a strong counter-cyclicality of the cross-sectional dispersion in subjective cash flow expectations.

7 Cyclicity of the Subjective User Cost of Labor

Overview The previous sections show that firms’ hiring decisions are heavily influenced by subjective cash flow expectations. This section examines whether expectations about the user cost of labor also contribute to hiring behavior, since it is a key component of the firm’s cash flows. Using survey data, I show that subjective wage expectations are significantly less cyclical than realized wages, implying that firms perceive labor costs as more rigid than they actually are.

To account for the possibility that wages depend on the economic conditions at the start of the job, I use survey expectations from the SCE to measure the user cost of labor under subjective expectations.¹⁴

In the search and matching model, the user cost of labor is the difference in the expected present value of wages between two firm-worker matches that are formed in two consecutive periods. Existing work assumes full information rational expectations and show that this user cost is more cyclical than flow wages, as workers hired in recessions earn lower wages both when hired and over time (Kudlyak, 2014; Bils et al., 2023). This section relaxes that assumption by using survey-based measures of subjective wage expectations. If firms and workers perceive the future path of wages as rigid, the subjective user cost of labor may remain high even during recessions, dampening hiring and amplifying unemployment fluctuations.

Time-series evidence Figure 3 compares realized real wage growth with 1-year-ahead subjective wage growth forecasts from three sources: the Livingston Survey, the CFO Survey, and the Survey of Consumer Expectations (SCE). Actual wage growth is clearly cyclical, with declines during downturns and strong rebounds during recoveries. In contrast, subjective wage forecasts are far more stable over time. Even during major shocks, such as the 2008 financial crisis and the COVID-19 recession, survey respondents anticipated only modest wage adjustments. Forecast errors are persistent and systematically biased: wage growth forecasts overestimate during downturns and underestimate during expansions.

To formally assess the cyclicity of real wage growth, Table 4 panel (a) compares the relationship between changes in the unemployment rate and real wage growth across rational and subjective expectations. As a rational expectations benchmark, I use historical data on actual real wage growth to estimate the following regression, replicating existing estimates in the literature (e.g., Bils, 1985; Solon et al., 1994; Gertler et al., 2020):

$$\Delta \log w_t = \beta_0 + \beta_1 \Delta u_t + \varepsilon_t \quad (34)$$

where $\Delta \log w_t$ represents the actual annual log growth rate of real wages, Δu_t is the annual change in the unemployment rate, and ε_t is the error term. β_1 is the coefficient of interest and captures the cyclicity of real wage growth.


Under subjective expectations, I use survey data on expected real wage growth to estimate:

$$F_{t-1}[\Delta \log w_t] = \beta_0 + \beta_1 F_{t-1}[\Delta u_t] + \varepsilon_t \quad (35)$$

where $F_{t-1}[\Delta \log w_t]$ is the median survey forecast for the annual log growth rate of real wages, where the surveys are either from Livingston, CFO, or SCE. $F_{t-1}[\Delta u_t]$ is the median survey forecast of the annual change in the unemployment rate from the Survey of Professional Forecasters (SPF). The coefficient of interest β_1 measures the cyclicity of expected real wage growth as perceived by survey respondents.

¹⁴See Section D for more details about its measurement.

Figure 3: Real Wage Growth: Actual vs. Subjective Expectations



Notes: This figure plots ex-post realized outcomes (Actual) and 1-year ahead subjective expectations (Survey) of real wage growth. x axis denotes the date on which actual values were realized and the period on which the survey forecast is made, making the vertical distance between the actual and survey lines the forecast error. Subjective expectations \mathbb{F}_t are based on survey forecasts. Left panel compares actual values of annual log real wage growth against the median consensus forecasts from the Livingston survey, where wages are measured using average weekly earnings of production and nonsupervisory employees, manufacturing (CES3000000030). Right panel compares annual log real wage growth against median consensus forecasts from the CFO survey and the subjective user cost of labor measured from the Survey of Consumer Expectations (SCE), where wages are measured using average hourly earnings of production and nonsupervisory employees, total private (CEU0500000008). Actual values are deflated using the Consumer Price Index (CPIAUCSL). Livingston, CFO, and SCE survey expectations of nominal wage growth are deflated using median consensus forecasts of CPI inflation from the Livingston, SPF, and SCE surveys, respectively. The sample period for Livingston is semi-annual spanning 1961S1 to 2022S2, CFO survey is quarterly spanning 2001Q4 to 2022Q4, SCE is monthly spanning 2015M5 to 2022M12. NBER recessions are shown with gray shaded bars.

Table 4 panel (a) reports the estimates. Under rational expectations, actual real wage growth is clearly cyclical since it is significantly negatively related to changes in unemployment rates. The magnitude of the estimate is also consistent with prior estimates in the literature, with elasticities ranging from -3.05 to -3.46 depending on the sample period (Solon et al., 1994). In contrast, subjective wage growth expectations are acyclical, with small and statistically insignificant coefficients across all survey sources and sample periods. Notably, the magnitude of the estimated elasticity is an order of magnitude smaller, ranging from -0.20 to -0.97 depending on the survey measure and sample period.

Cross-Sectional evidence To explore these patterns at the individual level, I use microdata from the SCE to estimate subjective wage cyclicalities separately for new hires and incumbents. The regression specification relaxes the rational expectations assumption from Gertler et al. (2020) and includes an interaction between expected unemployment growth and the probability of being a new hire:

$$\mathbb{F}_{t-1}[\Delta \log w_{i,t}] = \beta_0 + \beta_1 \mathbb{F}_{t-1}[\Delta u_t] + \mathbb{F}_{t-1}[\mathbb{I}\{N_{i,t} = 1\}] \cdot [\beta_2 + \beta_3 \mathbb{F}_{t-1}[\Delta u_t]] + \varepsilon_{i,t} \quad (36)$$

where $\mathbb{F}_{t-1}[\Delta \log w_{i,t}]$ represents the time $t - 1$ subjective expectation of wage growth for worker i at time t . $\mathbb{F}_{t-1}[\Delta u_t]$ is the survey-based expectation of aggregate unemployment growth. The

Table 4: Cyclicalities of Real Wage Growth: Actual vs. Subjective Expectations

(a) Aggregate Time-Series						
Actual: $\Delta \log w_t = \beta_0 + \beta_1 \Delta u_t + \varepsilon_t$						
Subjective: $\mathbb{F}_{t-1}[\Delta \log w_t] = \beta_0 + \beta_1 \mathbb{F}_{t-1}[\Delta u_t] + \varepsilon_t$						
	1961S1-2022S2		2001Q4-2022Q4		2015M5-2022M12	
	Actual	Survey Median (Liv)	Actual	Survey Median (CFO)	Actual	Survey User Cost (SCE)
	(1)	(2)	(3)	(4)	(5)	(6)
Unemployment Rate	-0.0340***	-0.0020	-0.0305***	0.0006	-0.0346***	-0.0086
<i>t</i> -stat	(-3.8684)	(-0.1568)	(-4.2477)	(0.0800)	(-6.6994)	(-1.6332)
Adj. R^2	0.1021	0.0003	0.2557	0.0001	0.4719	0.0498
<i>N</i>	124	124	85	85	92	92
Frequency	SA	SA	Q	Q	M	M
Sector	Mfg	Mfg	Pvt	Pvt	Pvt	Pvt

(b) Worker-Level New Hire Effect		
Subjective: $\mathbb{F}_{t-1}[\Delta \log w_{i,t}] = \beta_0 + \beta_1 \mathbb{F}_{t-1}[\Delta u_t] + \mathbb{F}_{t-1}[\mathbb{I}\{N_{i,t} = 1\}] \cdot [\beta_2 + \beta_3 \mathbb{F}_{t-1}[\Delta u_t]] + \varepsilon_{i,t}$		
	2015M5-2022M12	
	Survey (SCE)	Survey (SCE)
	(1)	(2)
	First Difference	Fixed Effects
Unemployment Rate	-0.0048 (0.0029)	-0.0028 (0.0026)
New Hire	0.0036*** (0.0009)	0.0003 (0.0013)
Unemployment Rate \times New Hire	-0.0026 (0.0020)	-0.0059 (0.0035)
Adj. R^2	0.0011	0.0036
<i>N</i>	39,832	39,832
Frequency	M	M
Sector	Pvt	Pvt

Notes: Table reports estimates from time-series and worker-level regressions of annual log real wage growth on unemployment growth. Subjective expectations \mathbb{F}_t are based on survey forecasts. Panel (a) reports estimates from time-series regressions using the aggregate series. Panel (a) Columns (1)-(2) compare actual values of annual log real wage growth against the median consensus forecasts from the Livingston survey, where wages are measured using average weekly earnings of production and nonsupervisory employees, manufacturing (CES3000000030). Panel (a) Columns (3)-(6) compare actual values of annual log real wage growth against median consensus forecasts from the CFO survey and the subjective user cost of labor measured from the Survey of Consumer Expectations (SCE), where wages are measured using average hourly earnings of production and nonsupervisory employees, total private (CEU0500000008). Panel (b) reports worker-level estimates from regressions of SCE survey expectations of wage growth on survey expectations of unemployment growth, an indicator of whether the worker is a new hire, and the interaction between the two. Actual wage growth is deflated using the Consumer Price Index (CPIAUCSL). Livingston, CFO, and SCE survey expectations of nominal wage growth are deflated using median consensus forecasts of CPI inflation from the Livingston, SPF, and SCE surveys, respectively. Subjective expectations of unemployment rates are from 1-year ahead consensus median forecasts from the SPF. The sample period for Livingston is semi-annual spanning 1961S1 to 2022S2, CFO survey is quarterly spanning 2001Q4 to 2022Q4, SCE is monthly spanning 2015M5 to 2022M12. Panel (a): Newey-West corrected *t*-statistics with lags 2 (semi-annual), 4 (quarterly), 12 (monthly) are reported in parentheses; Panel (b): Standard errors clustered by worker are reported in parentheses. *sig. at 10%. **sig. at 5%. ***sig. at 1%.

indicator variable $\mathbb{I}\{N_{i,t} = 1\}$ equals one if the worker is newly hired and zero otherwise. Its expectation $\mathbb{E}_{t-1}[\mathbb{I}\{N_{i,t} = 1\}]$ is thus the subjective perceived probability that the worker will be newly hired next period. The interaction term $\mathbb{E}_{t-1}[\mathbb{I}\{N_{i,t} = 1\}] \cdot \mathbb{E}_{t-1}[\Delta u_t]$ captures the differential sensitivity of expected wage growth to unemployment changes for new hires relative to incumbents. The error term $\varepsilon_{i,t}$ accounts for individual-level deviations in expectations.

The coefficient β_1 captures the overall cyclicalities of subjective wage expectations, reflecting how much workers expect wages to change in response to shifts in aggregate unemployment. The coefficient β_2 measures the baseline difference in expected wage growth between new hires and existing workers. The interaction term β_3 determines whether new hires expect wages to be more sensitive to unemployment fluctuations than incumbents do.

The results in Table 4 panel (b) column (1) show that, even after controlling for differences between job stayers and new hires, subjective wage expectations are highly rigid and exhibit weak cyclicalities. The coefficient β_1 is negative but small, confirming the aggregate result in panel (a) that workers that are not new hires expect only mild wage adjustments in response to unemployment fluctuations. The estimate for β_2 is positive, suggesting that, on average, new hires expect higher wage growth than job stayers. The interaction term β_3 is negative but small in magnitude, implying that new hires do not expect substantially greater cyclicalities in wages compared to incumbents. Column (2) extends column (1) by including worker fixed effects to find similar results. These findings extend the results from aggregate regressions by showing that subjective wage expectations are highly rigid even at the individual level, regardless of job transitions. Both new hires and incumbents perceive only weak cyclical variation in wages.

Implications for macroeconomic models These findings could have important implications for macroeconomic models of unemployment fluctuations. If firms do not expect wages to fall during downturns, then the subjective user cost of labor remains high even as demand declines, suppressing job creation. This mechanism is consistent with models that rely on wage rigidity to explain labor market volatility (Shimer, 2005; Hall, 2005; Christiano et al., 2016). These results suggest that it could be reasonable for macroeconomists to introduce rigid wages under subjective expectations to explain the volatility of business cycle fluctuations.

Moreover, the persistence of subjective wage expectations may reflect underlying frictions in information processing. Survey data on wage expectations can help distinguish between alternative theories of wage formation. Unlike rational models where the timing of wage payments is irrelevant (Barro, 1977), models with sticky or inattentive expectations—such as those in Mankiw and Reis (2002) or Coibion and Gorodnichenko (2015)—can be better suited to capture the persistent behavior of expected wages.

Finally, the finding that subjective cost of labor is rigid suggests that volatile subjective cash flow expectations are unlikely to be driven by fluctuations in the user cost of labor. Instead, firms may be over-reacting to other components of profitability, such as revenue expectations or perceived demand conditions, rather than expected changes in labor costs.

8 Robustness Checks and Extensions

This section presents additional results that reinforce the main finding. Across multiple robustness checks, the evidence consistently shows that firms overweight expected cash flows and underweight discount rates under subjective expectations.

Model vs. Data Model simulations suggest that the large role of subjective cash flow news poses a quantitative challenge for existing search models formulated under rational expectations, which have emphasized time-varying discount rates to match the volatility of unemployment fluctuations. Table 5 compares the empirical variance decomposition from Table 2 with those implied by benchmark search-and-matching models. For each model, I set a sample length of 20 years, produce 1,000 simulations, estimate a variance decomposition of the job filling rate according to equation (13), and report the average across the simulated runs. All parameter values in the calibration use estimates from the original papers.

The Diamond-Mortensen-Pisarides and Hall (2017) models predict that discount rate fluctuations should explain more than 75% of the variance of job filling rates. The Kehoe et al. (2022) (KLMP) model predicts a more balanced decomposition, attributing 54% to discount rates and 32% to cash flows, consistent with the model’s amplification mechanism based on human capital accumulation. In contrast, the empirical decomposition using survey data shows that subjective expectations assign just 1.6% of the variation to discount rates, and 96.4% to cash flows. These results highlight a discrepancy between the predictions of rational models and observed survey expectations, underscoring the importance of belief distortions.

Decreasing Returns to Scale and Compositional Effects Stock market valuations reflect average profits, while hiring decisions depend on marginal profits (Borovickova and Borovička, 2017). Decreasing returns to scale can amplify unemployment fluctuations even under a rational framework by making the marginal value of hiring more sensitive to productivity shocks, prompting firms to adjust vacancies more aggressively in response (Elsby and Michaels, 2013; Kaas and Kircher, 2015). Allowing for decreasing returns to scale introduces the notion of firm size. Changes in the equilibrium firm size distribution over the business cycle can introduce a compositional effect that also contributes to fluctuations in the job filling rate.

This section relaxes the constant returns to scale (CRS) assumption by allowing for decreasing returns to scale (DRS) in the production function. Assume that $Y_t = F(L_t) = A_t L_t^\alpha$, where A_t is an exogenous productivity process with $0 < \alpha < 1$. This introduces a “DRS wedge” between marginal and average profits:

$$\pi_t L_t - \kappa V_t = \alpha A_t L_t^\alpha - W_t L_t - \kappa V_t = E_t - (1 - \alpha) Y_t \quad (37)$$

where $E_t \equiv \Pi_t - \kappa V_t$ is the firm’s earnings, $\Pi_t \equiv Y_t - W_t L_t = A_t L_t^\alpha - W_t L_t$ is the total profit before wages $W_t L_t$ and vacancy posting costs κV_t , and $\pi_t = \frac{\partial \Pi_t}{\partial L_t}$ is the marginal profit from

Table 5: Time-Series Decomposition of the Job Filling Rate: Model vs. Data

Dep. Var. Horizon h (Years)	Discount Rate 5	(-) Cash Flow 5	(-) Price-Earnings 5	Residual 5
Subjective Expectations $\log q_t = c_q + \mathbb{F}_t[r_{t,t+h}] - \mathbb{F}_t[e_{t,t+h}] - \mathbb{F}_t[pe_{t,t+h}]$				
Data (Survey) t -stat	0.016 (0.180)	0.964*** (7.016)	0.027 (0.077)	-0.007
Rational Expectations $\log q_t = c_q + \mathbb{E}_t[r_{t,t+h}] - \mathbb{E}_t[e_{t,t+h}] - \mathbb{E}_t[pe_{t,t+h}]$				
Data (Machine) t -stat	0.511*** (2.765)	0.181*** (2.651)	0.291 (1.249)	0.017
Model (DMP) t -stat	0.782*** (12.334)	0.017** (1.992)	0.201*** (47.883)	0.000
Model (Hall) t -stat	0.838*** (12.000)	0.073 (1.387)	0.088 (1.074)	0.000
Model (KLMP) t -stat	0.543*** (4.484)	0.319 (0.937)	0.138*** (16.392)	0.000

Notes: Table compares the variance decomposition estimated from the data (Table 2) against the implied decomposition from simulations of alternative search-and-matching models. The models are simulated over 20 year samples, and all parameter values in the calibration use estimates from the original papers. DMP: Diamond-Mortensen-Pissarides Model; Hall: Hall (2017); KLMP: Kehoe et al. (2022). Subjective expectations \mathbb{F}_t are based on survey forecasts of CFOs and IBES financial analysts. Rational expectations \mathbb{E}_t are based on machine learning forecasts from Long Short-Term Memory (LSTM) neural networks. Newey-West corrected t -statistics with lags = 4 are reported in parentheses: *sig. at 10%. **sig. at 5%. ***sig. at 1%.

hiring. The second term $(1 - \alpha)Y_t$ is a “DRS wedge” that captures the gap between the average profit and marginal profit. Under DRS, the firm’s hiring condition becomes:

$$\frac{\kappa}{q_t} = \mathbb{F}_t \left[\sum_{j=1}^{\infty} \frac{1}{R_{t,t+j}} \left(\frac{E_{t+j}}{L_{t+1}} - (1 - \alpha) \frac{Y_{t+j}}{L_{t+1}} \right) \right] \quad (38)$$

Express aggregate earning-employment and output-employment ratios as the employment-weighted average of firm-level ratios

$$\frac{\kappa}{q_t} = \mathbb{F}_t \left[\sum_i \sum_{j=1}^{\infty} \frac{1}{R_{t,t+j}} \left(\frac{E_{i,t+j}}{L_{i,t+1}} - (1 - \alpha) \frac{Y_{i,t+j}}{L_{i,t+1}} \right) \frac{L_{i,t+1}}{L_{t+1}} \right] \quad (39)$$

Define $S_{i,t+1} \equiv \frac{L_{i,t+1}}{L_{t+1}}$ as the employment share, $EL_{i,t+j} \equiv E_{i,t+j}/L_{i,t+1}$ the earnings-employment ratio, and $YL_{i,t+j} \equiv Y_{i,t+j}/L_{i,t+1}$ the output-employment ratio of firm i . Log linearize the expression around the steady state

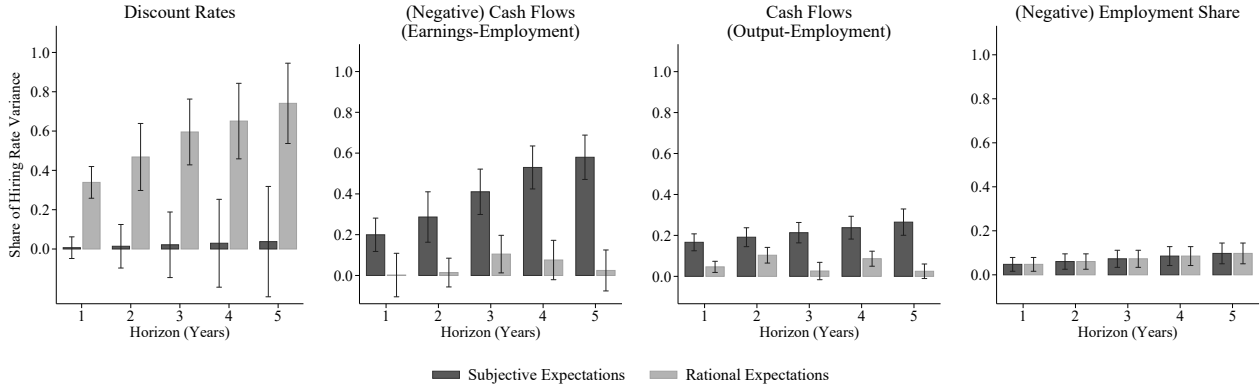
$$\log q_t = \sum_{j=1}^{\infty} \sum_i \left[\underbrace{\mathbb{F}_t[\rho_{r,i,j} r_{t,t+j}]}_{\text{Discount Rate}} - \underbrace{\mathbb{F}_t[\rho_{el,i,j} el_{i,t+j}]}_{\text{Cash Flow (Earnings-Employment)}} + \underbrace{\mathbb{F}_t[\rho_{yl,i,j} yl_{i,t+j}]}_{\text{Cash Flow (Output-Employment)}} - \underbrace{\mathbb{F}_t[\rho_{s,i,j} s_{i,t+1}]}_{\text{Employment Share}} \right] \quad (40)$$

where $r_{t,t+j}$, $el_{i,t+j}$, $yl_{i,t+j}$, and $s_{i,t+1}$ denote log deviations of $R_{t,t+j}$, $EL_{i,t+j}$, $YL_{i,t+j}$, and $S_{i,t+1}$ from the steady state state, respectively. The coefficients $\rho_{r,i,j} = \rho_{s,i,j} \equiv \frac{\bar{q}}{\kappa} \frac{(\bar{EL}_i + (1-\alpha)\bar{Y}\bar{L}_i) \cdot \bar{S}_i}{\bar{R}^j}$,

$\rho_{el,i,j} \equiv \frac{\bar{q}}{\kappa} \frac{\overline{EL}_i \bar{S}_i}{\bar{R}^j}$, and $\rho_{yl,i,j} \equiv (1 - \alpha) \frac{\bar{q}}{\kappa} \frac{\overline{YL}_i \bar{S}_i}{\bar{R}^j}$ are functions of steady-state values and linearization constants. $\alpha = 0.72$ comes from the labor share, $\kappa = 0.133$ comes from the flow vacancy cost (Elsby and Michaels, 2013). $\bar{q} = 0.631$, $\bar{R} = 1.04$, $\overline{EL} = 0.014$, $\overline{YL} = 0.074$ are long-run sample averages. Finally, approximate the infinite sum by truncating up to h periods.

The expected output-employment ratio $\mathbb{F}_t[y_{l,i,t+j}]$ captures the DRS wedge, and the employment share $s_{i,t+1}$ captures compositional effects of changes in the firm size distribution. I measure the expected output-employment ratio $\mathbb{F}_t[y_{l,i,t+j}]$ by using IBES sales forecasts. Figure 4 shows that under subjective expectations, the output-employment term accounts for roughly 40% of the variation in the job filling rate, while the earnings-employment term explains slightly less than 60%. The compositional term is small. These results confirm that even under DRS, subjective cash flow expectations—whether expressed in average or marginal terms—remain the dominant driver of hiring fluctuations.

Figure 4: Time-Series Decomposition of the Job Filling Rate: Decreasing Returns to Scale



Notes: This figure illustrates the components of the time-series decomposition of aggregate job filling rate under decreasing returns to scale, based on equation (40). The components of the decomposition are expected present discounted values of discount rate, earnings-employment ratio, output-employment ratio, and the employment share. The light bars show the contributions to the job filling rate obtained under rational expectations. The dark bars show the contributions to the time-series variation in the job filling rate obtained in subjective expectations. Subjective expectations \mathbb{F}_t are based on survey forecasts of CFOs and IBES financial analysts. Rational expectations \mathbb{E}_t are based on machine learning forecasts from Long Short-Term Memory (LSTM) neural networks. The sample is quarterly from 2005Q1 to 2021Q4. Each bar shows Newey-West 95% confidence intervals with lags = 4.

Predictability of Survey Forecast Errors and Bias To assess whether survey expectations systematically deviate from rational expectations, panel (a) of Table 6 estimates Coibion and Gorodnichenko (2015) regressions of forecast errors and bias on forecast revisions:

$$y_{t,t+h} - \mathbb{F}_t[y_{t,t+h}] = \beta_{FE,0} + \beta_{FE,1}[\mathbb{F}_t[y_{t,t+h}] - \mathbb{F}_{t-1}[y_{t,t+h}]] + \beta_{FE,2}\mathbb{F}_{t-1}[y_{t,t+h}] + \varepsilon_{FE,t} \quad (41)$$

The negative estimate on the forecast revision ($\beta_{FE,1} < 0$) confirms that an upward revision in the forecast is an over-reaction that in turn predicts a negative forecast error, in which the survey forecast is above its ex-post realized value. The over-reaction persists for multiple periods until the upward revision eventually predicts a negative forecast error.

To assess how much of the over-reaction from panel (a) is driven by systematic ex-ante biases in survey expectations, panel (b) of Table 6 estimates regressions of the survey bias on the survey forecast revisions:

$$\mathbb{E}_t[y_{t,t+h}] - \mathbb{F}_t[y_{t,t+h}] = \beta_{B,0} + \beta_{B,1}[\mathbb{F}_t[y_{t,t+h}] - \mathbb{F}_{t-1}[y_{t,t+h}]] + \beta_{B,2}\mathbb{F}_{t-1}[y_{t,t+h}] + \varepsilon_{B,t} \quad (42)$$

The negative estimate on the forecast revision ($\beta_{B,1} < 0$) suggests that the bias component of the survey expectation over-reacts to news. The estimates become increasingly negative with the forecast horizon, suggesting that the over-reaction is stronger for survey expectations over longer horizons. The substantial size of the estimates compared to panel (a) highlights that a substantial share of the over-reaction can be attributed to biases in survey expectations that could have been identified ex-ante. At the 5 year horizon, about $-0.735 / -1.014 = 72.4\%$ of the over-reaction identified from survey forecast errors can be attributed to ex-ante biases in survey expectations.

Table 6: Predictability of Survey Forecast Errors and Bias

Horizon h (Years)	1	2	3	4	5
(a) Survey Forecast Errors					
$y_{t,t+h} - \mathbb{F}_t[y_{t,t+h}] = \beta_{FE,0} + \beta_{FE,1}[\mathbb{F}_t[y_{t,t+h}] - \mathbb{F}_{t-1}[y_{t,t+h}]] + \beta_{FE,2}\mathbb{F}_{t-1}[y_{t,t+h}] + \varepsilon_{FE,t}$					
Discount Rate	-1.473	-2.211**	-2.257***	-2.638***	-3.191***
t -stat	(-1.020)	(-2.476)	(-2.730)	(-4.448)	(-6.020)
Cash Flow	-0.550**	-0.685***	-0.942***	-0.916***	-1.014***
t -stat	(-2.461)	(-4.745)	(-9.548)	(-8.119)	(-7.380)
Price-Earnings	-0.574***	-0.006	-0.313**	-0.475***	-0.885***
t -stat	(-2.924)	(-0.034)	(-2.348)	(-4.314)	(-6.154)
N	68	68	68	68	68
(b) Survey Bias					
$\mathbb{E}_t[y_{t,t+h}] - \mathbb{F}_t[y_{t,t+h}] = \beta_{B,0} + \beta_{B,1}[\mathbb{F}_t[y_{t,t+h}] - \mathbb{F}_{t-1}[y_{t,t+h}]] + \beta_{B,2}\mathbb{F}_{t-1}[y_{t,t+h}] + \varepsilon_{B,t}$					
Discount Rate	-1.521***	-1.187***	-1.373***	-0.958***	-1.065***
t -stat	(-4.470)	(-2.997)	(-2.930)	(-2.849)	(-3.705)
Cash Flow	-0.227	-0.302***	-0.632***	-0.641***	-0.735***
t -stat	(-1.599)	(-2.641)	(-6.438)	(-5.659)	(-7.077)
Price-Earnings	-0.547***	-0.273***	-0.739***	-0.503***	-0.858***
t -stat	(-5.813)	(-3.928)	(-2.614)	(-2.614)	(-6.835)
N	68	68	68	68	68

Notes: Table reports slope coefficients $\beta_{FE,1}$ and $\beta_{B,1}$ from regressions of the survey forecast error and its bias on the survey forecast revisions. $y_{t,t+h}$ denotes the variable y to be predicted h years ahead of time t : h year present discounted values of discount rates ($r_{t,t+h}$), cash flows ($e_{t,t+h}$), and log price-earnings ratios ($pe_{t,t+h}$). Subjective expectations \mathbb{F}_t are based on survey forecasts of CFOs and IBES financial analysts. Rational expectations \mathbb{E}_t are based on machine learning forecasts from Long Short-Term Memory (LSTM) neural networks. The survey bias is defined as ex-ante deviations of the survey forecast from its machine learning benchmark: $\mathbb{E}_t[y_{t,t+h}] - \mathbb{F}_t[y_{t,t+h}]$. The sample is quarterly from 2005Q1 to 2021Q4. Newey-West corrected t -statistics with lags = 4 are reported in parentheses: *sig. at 10%. **sig. at 5%. ***sig. at 1%.

Additional Results Appendix Section A presents the following supplementary results. Table A.1 reports summary statistics for each variable used in the empirical analysis. Variance decompositions based on a vector autoregression (VAR) (Table A.2) confirm that under subjective

expectations, cash flows account for nearly all variation in job filling rates, while under rational expectations, discount rates dominate. The terminal value captured by future price-earnings ratio diminishes to near zero by the 5-year horizon, consistent with the main results.

The results are robust to alternative specifications, including first-differences (Table A.3), conditioning on control variables such as lagged job filling rates and survey forecasts (Table A.4), and alternative survey sources such as the Bloomberg and CFO surveys (Tables A.5, A.6, and A.7). Cross-sectional results continue to hold when constructing Bartik instruments for subjective discount rate and cash flow expectations (Table A.8).

The following results validate the search model from Section 2 to ensure that it provides a reasonable description of labor market fluctuations. Job filling rates are closely linked to forward-looking price-employment (Table A.9) and price-earnings ratios (Table A.10). Estimating a similar decomposition for labor market tightness assuming a constant returns to scale matching function (Table A.11) shows that subjective earnings expectations account for nearly all variation, while rational models emphasize discount rates, echoing the findings for job filling rates.

The job filling rate strongly predicts survey biases, defined as the deviation of subjective expectations from a rational expectations benchmark (Table A.12). Predictable deviations from rational expectations can explain a substantial share of job filling rate variation. Decompositions under risk-neutral expectations show that risk-free rates do not play a significant role in explaining job filling rates (Tables A.13 and A.14). The risk-neutral estimates continue to show that subjective expectations to overweight long-horizon cash flows. Ex-post decompositions confirm that realized discount rate news—not cash flows—drive hiring outcomes, aligning with machine learning measures of rational expectations (Table A.15).

9 Conclusion

This paper examines how belief distortions affect unemployment fluctuations by comparing survey-based subjective expectations with machine learning forecasts that proxy for rational beliefs. Using a decomposition of the job filling rate grounded in a search-and-matching model, I show that firms’ hiring decisions under subjective expectations are driven almost entirely by expected future cash flows, with little role for discount rates. In contrast, rational expectations assign a dominant role to discount rate variation. Cross-sectional evidence using firm and state level data reinforces this mechanism: states where firms exhibit stronger distortions in cash flow expectations experience larger swings in job filling rates. Additionally, I find that subjective wage expectations are far less cyclical than observed wages, both for new hires and continuing workers. This perceived rigidity keeps the user cost of labor persistently elevated in recessions, amplifying unemployment fluctuations by reducing firms’ willingness to hire workers.

These findings suggest that incorporating subjective expectations into existing models can enhance our understanding of labor market dynamics. While standard rational models imply that agents correctly understand the relative importance of discount rates and cash flows, belief dis-

tortions provide a complementary channel through which non-rational beliefs can influence hiring and unemployment volatility. This study opens several directions for future research. Investigating how firms update their expectations based on their size, industry, or financial constraints could help refine the microfoundations for their belief-driven hiring behavior. The subjective beliefs disciplined by survey forecasts can also be embedded into quantitative search and matching models of the labor market, allowing for counterfactual simulations that can test how labor markets would behave under alternative belief formation mechanisms or policy interventions.

References

- Acemoglu, Daron and William B. Hawkins**, “Search with multi-worker firms,” *Theoretical Economics*, 2014, 9 (3), 583–628.
- Acharya, Sushant and Shu Lin Wee**, “Rational Inattention in Hiring Decisions,” *American Economic Journal: Macroeconomics*, January 2020, 12 (1), 1–40.
- Adam, Klaus, Dmitry Matveev, and Stefan Nagel**, “Do survey expectations of stock returns reflect risk adjustments?,” *Journal of Monetary Economics*, 2021, 117, 723–740.
- Afrouzi, Hassan, Spencer Y Kwon, Augustin Landier, Yueran Ma, and David Thesmar**, “Overreaction in Expectations: Evidence and Theory,” *The Quarterly Journal of Economics*, 03 2023, 138 (3), 1713–1764.
- Amior, Michael and Alan Manning**, “The Persistence of Local Joblessness,” *American Economic Review*, July 2018, 108 (7), 1942–70.
- Andreou, Elena, Eric Ghysels, and Andros Kourtellos**, “Should macroeconomic forecasters use daily financial data and how?,” *Journal of Business & Economic Statistics*, 2013, 31 (2), 240–251.
- Augenblick, Ned, Eben Lazarus, and Michael Thaler**, “Overinference from Weak Signals and Underinference from Strong Signals,” *The Quarterly Journal of Economics*, 10 2024, 140 (1), 335–401.
- Baker, Scott, Nicholas Bloom, Steven J Davis, and Marco Sammon**, “What triggers stock market jumps?,” 2019. Unpublilshed manuscript, Stanford University.
- Barberis, Nicholas, Andrei Shleifer, and Robert Vishny**, “A model of investor sentiment1We are grateful to the NSF for financial support, and to Oliver Blanchard, Alon Brav, John Campbell (a referee), John Cochrane, Edward Glaeser, J.B. Heaton, Danny Kahneman, David Laibson, Owen Lamont, Drazen Prelec, Jay Ritter (a referee), Ken Singleton, Dick Thaler, an anonymous referee, and the editor, Bill Schwert, for comments.1,” *Journal of Financial Economics*, 1998, 49 (3), 307–343.
- Barnichon, Regis**, “Building a composite Help-Wanted Index,” *Economics Letters*, 2010, 109 (3), 175–178.
- Barro, Robert J.**, “Long-term contracting, sticky prices, and monetary policy,” *Journal of Monetary Economics*, 1977, 3 (3), 305–316.
- Basu, S. and C.L. House**, “Allocative and Remitted Wages,” in J. B. Taylor and Harald Uhlig, eds., *Handbook of Macroeconomics*, Vol. 2 of *Handbook of Macroeconomics*, Elsevier, 2016, chapter 0, pp. 297–354.

- Bauer, Michael D. and Eric T. Swanson**, “An Alternative Explanation for the ”Fed Information Effect”,” *American Economic Review*, March 2023, *113* (3), 664–700.
- Belo, Frederico, Andres Donangelo, Xiaoji Lin, and Ding Luo**, “What Drives Firms’ Hiring Decisions? An Asset Pricing Perspective,” *The Review of Financial Studies*, 02 2023. hhad012.
- Ben-David, Itzhak, John R. Graham, and Campbell R. Harvey**, “Managerial Miscalibration,” *The Quarterly Journal of Economics*, 09 2013, *128* (4), 1547–1584.
- Beraja, Martin, Erik Hurst, and Juan Ospina**, “The Aggregate Implications of Regional Business Cycles,” *Econometrica*, 2019, *87* (6), 1789–1833.
- Bhandari, Anmol, Jaroslav Borovička, and Paul Ho**, “Survey Data and Subjective Beliefs in Business Cycle Models,” *The Review of Economic Studies*, 05 2024, p. rdae054.
- Bianchi, Francesco, Cosmin Ilut, and Hikaru Saijo**, “Diagnostic Business Cycles,” *The Review of Economic Studies*, 02 2023, p. rdad024.
- , **Cosmin L Ilut, and Hikaru Saijo**, “Smooth Diagnostic Expectations,” Working Paper 32152, National Bureau of Economic Research February 2024.
- , **Sydney C. Ludvigson, and Sai Ma**, “Belief Distortions and Macroeconomic Fluctuations,” *American Economic Review*, July 2022, *112* (7), 2269–2315.
- , **Sydney C Ludvigson, and Sai Ma**, “What Hundreds of Economic News Events Say About Belief Overreaction in the Stock Market,” Working Paper 32301, National Bureau of Economic Research 4 2024.
- Bils, Mark J.**, “Real Wages over the Business Cycle: Evidence from Panel Data,” *Journal of Political Economy*, 1985, *93* (4), 666–689.
- Bils, Mark, Marianna Kudlyak, and Paulo Lins**, “The Quality-Adjusted Cyclical Price of Labor,” *Journal of Labor Economics*, 2023, *41* (S1), S13–S59.
- Blanchard, Olivier and Jordi Galí**, “Labor Markets and Monetary Policy: A New Keynesian Model with Unemployment,” *American Economic Journal: Macroeconomics*, April 2010, *2* (2), 1–30.
- Blei, David M, Andrew Y Ng, and Michael I Jordan**, “Latent dirichlet allocation,” *Journal of machine Learning research*, 2003, *3* (Jan), 993–1022.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer**, “Diagnostic Expectations and Credit Cycles,” *The Journal of Finance*, 2018, *73* (1), 199–227.

- , – , – , and **Stephen J Terry**, “Real Credit Cycles,” Working Paper 28416, National Bureau of Economic Research 1 2021.
- , – , **Rafael La Porta**, and **Andrei Shleifer**, “Diagnostic Expectations and Stock Returns,” *The Journal of Finance*, 2019, *74* (6), 2839–2874.
- , – , **Rafael La Porta**, and **Andrei Shleifer**, “Belief Overreaction and Stock Market Puzzles,” *Journal of Political Economy*, 2024, *132* (5), 1450–1484.
- , – , – , **Matthew OBrien**, and **Andrei Shleifer**, “Long-Term Expectations and Aggregate Fluctuations,” *NBER Macroeconomics Annual*, 2024, *38*, 311–347.
- , – , **Yueran Ma**, and **Andrei Shleifer**, “Overreaction in Macroeconomic Expectations,” *American Economic Review*, September 2020, *110* (9), 2748–82.
- Borovickova, Katarina** and **Jaroslav Borovička**, “Discount Rates and Employment Fluctuations,” 2017 Meeting Papers 1428, Society for Economic Dynamics 2017.
- Bybee, Leland, Bryan T Kelly, Asaf Manela, and Dacheng Xiu**, “Business news and business cycles,” Technical Report, National Bureau of Economic Research 2021.
- , – , – , and – , “Business News and Business Cycles,” *The Journal of Finance*, 2024, *79* (5), 3105–3147.
- Campbell, John Y. and Robert J. Shiller**, “The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors,” *The Review of Financial Studies*, 1988, *1* (3), 195–228.
- Candia, Bernardo, Olivier Coibion, and Yuriy Gorodnichenko**, “Communication and the Beliefs of Economic Agents,” Working Paper 27800, National Bureau of Economic Research September 2020.
- Chen, Long, Zhi Da, and Xinlei Zhao**, “What Drives Stock Price Movements?,” *The Review of Financial Studies*, 02 2013, *26* (4), 841–876.
- Chodorow-Reich, Gabriel and Johannes Wieland**, “Secular Labor Reallocation and Business Cycles,” *Journal of Political Economy*, 2020, *128* (6), 2245–2287.
- Christiano, Lawrence J., Martin S. Eichenbaum, and Mathias Trabandt**, “Unemployment and Business Cycles,” *Econometrica*, 2016, *84* (4), 1523–1569.
- Cochrane, John H.**, “Production-Based Asset Pricing and the Link Between Stock Returns and Economic Fluctuations,” *The Journal of Finance*, 1991, *46* (1), 209–237.

- , “The Dog That Did Not Bark: A Defense of Return Predictability,” *The Review of Financial Studies*, 09 2007, *21* (4), 1533–1575.
- Cochrane, John H.**, “Macro-Finance,” *Review of Finance*, 03 2017, *21* (3), 945–985.
- Coibion, Olivier and Yuriy Gorodnichenko**, “Information rigidity and the expectations formation process: A simple framework and new facts,” *American Economic Review*, 2015, *105* (8), 2644–78.
- , – , and **Saten Kumar**, “How Do Firms Form Their Expectations? New Survey Evidence,” *American Economic Review*, September 2018, *108* (9), 2671–2713.
- Cong, Lin William, Ke Tang, Jingyuan Wang, and Yang Zhang**, “Deep Sequence Modeling: Development and Applications in Asset Pricing,” *The Journal of Financial Data Science*, December 2020, *3* (1), 28–42.
- Cooper, Rick A., Theodore E. Day, and Craig M. Lewis**, “Following the leader:: a study of individual analysts’ earnings forecasts,” *Journal of Financial Economics*, 2001, *61* (3), 383–416.
- David, Joel M., Lukas Schmid, and David Zeke**, “Risk-adjusted capital allocation and misallocation,” *Journal of Financial Economics*, 2022, *145* (3), 684–705.
- Diamond, Peter A.**, “Wage Determination and Efficiency in Search Equilibrium,” *The Review of Economic Studies*, 04 1982, *49* (2), 217–227.
- Elsby, Michael W. L. and Ryan Michaels**, “Marginal Jobs, Heterogeneous Firms, and Unemployment Flows,” *American Economic Journal: Macroeconomics*, January 2013, *5* (1), 1–48.
- Faberman, R. Jason, Andreas I. Mueller, Aysegül Şahin, and Giorgio Topa**, “Job Search Behavior Among the Employed and Non-Employed,” *Econometrica*, 2022, *90* (4), 1743–1779.
- Fujita, Shigeru and Garey Ramey**, “Exogenous versus Endogenous Separation,” *American Economic Journal: Macroeconomics*, May 2012, *4* (4), 68–93.
- Gabaix, Xavier**, “Behavioral inattention,” in B. Douglas Bernheim, Stefano DellaVigna, and David Laibson, eds., *Handbook of Behavioral Economics - Foundations and Applications 2*, Vol. 2 of *Handbook of Behavioral Economics: Applications and Foundations 1*, North-Holland, 2019, pp. 261–343.
- Gavazza, Alessandro, Simon Mongey, and Giovanni L. Violante**, “Aggregate Recruiting Intensity,” *American Economic Review*, August 2018, *108* (8), 2088–2127.

- Gennaioli, Nicola, Yueran Ma, and Andrei Shleifer**, “Expectations and Investment,” *NBER Macroeconomics Annual*, 2016, *30*, 379–431.
- Gertler, Mark and Antonella Trigari**, “Unemployment Fluctuations with Staggered Nash Wage Bargaining,” *Journal of Political Economy*, 2009, *117* (1), 38–86.
- , **Christopher Huckfeldt, and Antonella Trigari**, “Unemployment Fluctuations, Match Quality, and the Wage Cyclicalilty of New Hires,” *The Review of Economic Studies*, 02 2020, *87* (4), 1876–1914.
- Gormsen, Niels Joachim and Kilian Huber**, “Corporate Discount Rates,” Working Paper 31329, National Bureau of Economic Research 06 2023.
- and **Ralph S J Koijen**, “Coronavirus: Impact on Stock Prices and Growth Expectations,” *The Review of Asset Pricing Studies*, 09 2020, *10* (4), 574–597.
- Greenwood, Robin and Andrei Shleifer**, “Expectations of returns and expected returns,” *The Review of Financial Studies*, 2014, *27* (3), 714–746.
- Gu, Shihao, Bryan Kelly, and Dacheng Xiu**, “Empirical Asset Pricing via Machine Learning,” *The Review of Financial Studies*, 02 2020, *33* (5), 2223–2273.
- Hall, Robert E.**, “Employment Fluctuations with Equilibrium Wage Stickiness,” *American Economic Review*, March 2005, *95* (1), 50–65.
- , “High Discounts and High Unemployment,” *American Economic Review*, 2 2017, *107* (2), 305–30.
- Hansen, Lars Peter**, “Beliefs, Doubts and Learning: Valuing Macroeconomic Risk,” *American Economic Review*, May 2007, *97* (2), 1–30.
- , **John C. Heaton, and Nan Li**, “Intangible Risk,” in “Measuring Capital in the New Economy,” University of Chicago Press, October 2005.
- Hayashi, Fumio**, “Tobin’s Marginal q and Average q: A Neoclassical Interpretation,” *Econometrica*, 1982, *50* (1), 213–224.
- Hazell, Jonathon and Bledi Taska**, “Downward Rigidity in the Wage for New Hires,” IZA Discussion Papers 16512, Institute of Labor Economics (IZA) October 2023.
- , **Juan Herreño, Emi Nakamura, and Jón Steinsson**, “The Slope of the Phillips Curve: Evidence from U.S. States,” *The Quarterly Journal of Economics*, 02 2022, *137* (3), 1299–1344.

- J., Jr. Dennis William**, “Small Business Credit In a Deep Recession,” Report, NFIB Research Foundation, United States February 2010. Financial Crisis Inquiry Commission (FCIC) Case Series.
- Kaas, Leo and Philipp Kircher**, “Efficient Firm Dynamics in a Frictional Labor Market,” *American Economic Review*, October 2015, *105* (10), 3030–60.
- Kehoe, Patrick J, Pierlauro Lopez, Virgiliu Midrigan, and Elena Pastorino**, “Asset Prices and Unemployment Fluctuations: A Resolution of the Unemployment Volatility Puzzle,” *The Review of Economic Studies*, 08 2022, *90* (3), 1304–1357.
- Kehoe, Patrick J., Virgiliu Midrigan, and Elena Pastorino**, “Debt Constraints and Employment,” *Journal of Political Economy*, 2019, *127* (4), 1926–1991.
- Kilic, Mete and Jessica A Wachter**, “Risk, Unemployment, and the Stock Market: A Rare-Event-Based Explanation of Labor Market Volatility,” *The Review of Financial Studies*, 01 2018, *31* (12), 4762–4814.
- Kogan, Leonid and Dimitris Papanikolaou**, “Economic Activity of Firms and Asset Prices,” *Annual Review of Financial Economics*, 2012, *4* (Volume 4, 2012), 361–384.
- Korniotis, George M.**, “Habit Formation, Incomplete Markets, and the Significance of Regional Risk for Expected Returns,” *The Review of Financial Studies*, 08 2008, *21* (5), 2139–2172.
- Kothari, S.P., Eric So, and Rodrigo Verdi**, “Analysts’ Forecasts and Asset Pricing: A Survey,” *Annual Review of Financial Economics*, 2016, *8* (Volume 8, 2016), 197–219.
- Kudlyak, Marianna**, “The cyclical cost of the user cost of labor,” *Journal of Monetary Economics*, 2014, *68*, 53–67.
- Leduc, Sylvain and Zheng Liu**, “The Weak Job Recovery in a Macro Model of Search and Recruiting Intensity,” *American Economic Journal: Macroeconomics*, January 2020, *12* (1), 310–43.
- Lettau, Martin and Sydney Ludvigson**, “Time-varying risk premia and the cost of capital: An alternative implication of the Q theory of investment,” *Journal of Monetary Economics*, 2002, *49* (1), 31–66.
- Liu, Laura Xiaolei, Toni M. Whited, and Lu Zhang**, “Investment-Based Expected Stock Returns,” *Journal of Political Economy*, 2009, *117* (6), 1105–1139.
- Liu, Yukun**, “Labor-based asset pricing,” *SSRN*, 2021.

- Ljungqvist, Lars and Thomas J. Sargent**, “The Fundamental Surplus,” *American Economic Review*, September 2017, *107* (9), 2630–65.
- Ludvigson, Sydney C. and Serena Ng**, “The empirical risk-return relation: A factor analysis approach,” *Journal of Financial Economics*, January 2007, *83* (1), 171–222.
- Ma, Yueran, Tiziano Ropele, David Sraer, and David Thesmar**, “A Quantitative Analysis of Distortions in Managerial Forecasts,” Working Paper 26830, National Bureau of Economic Research March 2020.
- Mankiw, N. Gregory and Ricardo Reis**, “Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve,” *The Quarterly Journal of Economics*, 11 2002, *117* (4), 1295–1328.
- Manning, Alan and Barbara Petrongolo**, “How Local Are Labor Markets? Evidence from a Spatial Job Search Model,” *American Economic Review*, October 2017, *107* (10), 2877–2907.
- Meeuwis, Maarten, Dimitris Papanikolaou, Jonathan L Rothbaum, and Lawrence D.W. Schmidt**, “Time-Varying Risk Premia, Labor Market Dynamics, and Income Risk,” Working Paper 31968, National Bureau of Economic Research December 2023.
- Menzio, Guido**, “Stubborn Beliefs in Search Equilibrium,” *NBER Macroeconomics Annual*, 2023, *37*, 239–297.
- Mitra, Indrajit, Taeuk Seo, and Yu Xu**, “High Discounts and Low Fundamental Surplus: An Equivalence Result for Unemployment Fluctuations,” *Management Science*, 2024, *70* (6), 4051–4068.
- Mortensen, Dale T.**, “The Matching Process as a Noncooperative Bargaining Game,” in “The Economics of Information and Uncertainty” NBER Chapters, National Bureau of Economic Research, Inc, May 1982, pp. 233–258.
- Mueller, Andreas I., Johannes Spinnewijn, and Giorgio Topa**, “Job Seekers’ Perceptions and Employment Prospects: Heterogeneity, Duration Dependence, and Bias,” *American Economic Review*, January 2021, *111* (1), 324–63.
- Nagel, Stefan and Zhengyang Xu**, “Asset Pricing with Fading Memory,” *The Review of Financial Studies*, 08 2021, *35* (5), 2190–2245.
- and —, “Dynamics of Subjective Risk Premia,” Working Paper 29803, National Bureau of Economic Research 2 2022.
- O, Ricardo De La and Sean Myers**, “Subjective Cash Flow and Discount Rate Expectations,” *The Journal of Finance*, 2021, *76* (3), 1339–1387.

- , **Xiao Han**, and **Sean Myers**, “The Cross-section of Subjective Expectations: Understanding Prices and Anomalies,” *SSRN*, 2024.
- Oi, Walter Y. and Todd L. Idson**, “Firm size and wages,” *Handbook of Labor Economics*, 1999, 3, 2165–2214.
- Petrosky-Nadeau, Nicolas, Lu Zhang, and Lars-Alexander Kuehn**, “Endogenous Disasters,” *American Economic Review*, 8 2018, 108 (8), 2212–45.
- Pissarides, Christopher A.**, “The Unemployment Volatility Puzzle: Is Wage Stickiness the Answer?,” *Econometrica*, 2009, 77 (5), 1339–1369.
- Ropele, Tiziano, Yuriy Gorodnichenko, and Olivier Coibion**, “Inflation Expectations and Misallocation of Resources: Evidence from Italy,” *American Economic Review: Insights*, June 2024, 6 (2), 246–61.
- Sahin, Aysegul, Joseph Song, Giorgio Topa, and Giovanni L. Violante**, “Mismatch Unemployment,” *American Economic Review*, November 2014, 104 (11), 3529–64.
- Shimer, Robert**, “The Cyclical Behavior of Equilibrium Unemployment and Vacancies,” *American Economic Review*, 3 2005, 95 (1), 25–49.
- , “Reassessing the ins and outs of unemployment,” *Review of Economic Dynamics*, 2012, 15 (2), 127–148.
- Sims, Christopher A.**, “Implications of rational inattention,” *Journal of Monetary Economics*, 2003, 50 (3), 665–690. Swiss National Bank/Study Center Gerzensee Conference on Monetary Policy under Incomplete Information.
- Solon, Gary, Robert Barsky, and Jonathan A. Parker**, “Measuring the Cyclicalities of Real Wages: How Important is Composition Bias,” *The Quarterly Journal of Economics*, 1994, 109 (1), 1–25.
- Woodford, Michael**, “Imperfect Common Knowledge and the Effects of Monetary Policy,” Working Paper 8673, National Bureau of Economic Research December 2001.

A Appendix: Additional Results

A.1 Additional Stylized Facts

Table A.1: Summary statistics

Panel (a): Aggregate U.S.								
	Obs	Mean	St. Dev.	Min	p25	Median	p75	Max
$r_{t,t+5}$	68	0.351	0.321	-0.221	0.095	0.379	0.653	0.906
$\mathbb{F}_t[r_{t,t+5}]$	68	0.194	0.073	0.062	0.149	0.183	0.216	0.371
$\mathbb{E}_t[r_{t,t+5}]$	68	0.301	0.129	-0.003	0.234	0.315	0.374	0.661
$e_{t,t+5}$	68	1.389	0.301	0.032	1.362	1.460	1.543	1.911
$\mathbb{F}_t[e_{t,t+5}]$	68	1.872	0.134	1.593	1.767	1.901	2.000	2.068
$\mathbb{E}_t[e_{t,t+5}]$	68	1.399	0.198	1.074	1.231	1.399	1.587	1.765
pe_{t+5}	68	2.227	0.314	1.784	1.995	2.240	2.347	3.294
$\mathbb{F}_t[pe_{t,t+5}]$	68	1.683	0.228	1.286	1.515	1.685	1.763	2.228
$\mathbb{E}_t[pe_{t,t+5}]$	68	2.273	0.279	1.686	2.089	2.197	2.444	2.916
q_t	68	0.616	0.228	0.211	0.436	0.603	0.738	1.202
Panel (b): State-Level								
	Obs	Mean	St. Dev.	Min	p25	Median	p75	Max
$q_{i,t}$	16,302	0.141	0.219	-0.882	0.032	0.141	0.271	1.153
$\tilde{q}_{i,t}$	16,302	0.167	0.216	-0.346	0.062	0.167	0.298	0.681
Panel (c): Firm-Level								
	Obs	Mean	St. Dev.	Min	p25	Median	p75	Max
$\mathbb{F}_t[r_{f,i,t,t+5}]$	611,636	0.262	0.312	-0.477	0.091	0.185	0.330	3.399
$\mathbb{F}_t[e_{f,i,t,t+5}]$	611,636	0.078	2.196	-8.561	-1.366	-0.014	1.505	7.410
$\mathbb{F}_t[pe_{f,i,t,t+5}]$	611,636	2.583	0.715	-2.263	2.262	2.531	2.832	8.204
$\Delta\mathbb{F}_t[r_{f,i,t,t+5}]$	549,513	-0.010	0.233	-1.531	-0.115	-0.006	0.100	1.460
$\Delta\mathbb{F}_t[e_{f,i,t,t+5}]$	549,513	0.014	0.458	-2.668	-0.162	0.007	0.190	2.589
$\Delta\mathbb{F}_t[pe_{f,i,t,t+5}]$	549,513	-0.004	0.405	-2.286	-0.171	0.002	0.170	2.304

Notes: This table reports summary statistics for ex-post realized outcomes (Actual), subjective expectations (Survey), and machine expectations (Machine) of key dependent variables used in the variance decomposition. Panel (a) reports aggregate U.S. statistics, Panel (b) shows state-level statistics, and Panel (c) presents firm-level statistics. The dependent variables are $h = 5$ year present discounted values of discount rates $r_{t,t+h}$, cash flows $e_{t,t+h}$, and price-earnings ratios $pe_{t,t+h}$, as defined in equation (11). Aggregate labor market variables include the job filling rate q_t . State-level job filling rate is measured in two ways: (1) a direct measure calculated as the ratio of JOLTS total hires to job vacancies $q_{i,t}$, and (2) a Bartik shift-share instrument $\tilde{q}_{i,t}$ that interacts state-level employment shares in 2-digit NAICS industries with national industry-specific job filling rate fluctuations. Subjective expectations \mathbb{F}_t are based on survey forecasts from the CFO survey for stock returns, and IBES for earnings growth. Machine expectations are based on machine learning forecasts \mathbb{E}_t from Long Short-Term Memory (LSTM) neural networks $G(\mathcal{X}_t, \beta_{h,t})$, whose parameters $\beta_{h,t}$ are estimated in real time using \mathcal{X}_t , a large scale dataset of macroeconomic, financial, and textual data. The out-of-sample testing period for aggregate-level survey and machine forecasts is quarterly and spans 2005Q1 to 2021Q4. The sample for state and firm-level data are monthly and spans 2001M1 to 2021M12.

A.2 Variance Decomposition: VAR Estimates

To further validate the robustness of the variance decomposition in Table 2, I estimate a Vector Autoregression (VAR) model of the job filling rate and its key determinants under both subjective and rational expectations. The estimated system follows a VAR(1) specification:

$$X_{t+1} = AX_t + \varepsilon_{t+1} \quad (\text{A.1})$$

where the state vector X_t consists of the present discounted values of discount rates, cash flows, log price-earnings ratios, and the log job filling rate.

Table A.2 reports the estimates. Under rational expectations, discount rate fluctuations explain a moderate but significant share of the variance in job filling rates, rising from 36% at the one-year horizon to 62% at the 5-year horizon. Under subjective expectations, cash flow expectations overwhelmingly dominate, accounting for 59% at the one-year horizon and rising to 96% at the 5-year horizon. Subjective discount rate fluctuations play an insignificant role, remaining close to zero across all horizons. These findings confirm the main results in Table 2: Managers using subjective expectations place disproportionate weight on cash flow expectations, whereas models with rational expectations emphasize the role of discount rate fluctuations.

Table A.2: Variance Decomposition of Job Filling Rate: VAR(1)

Horizon h (Years)	1	2	3	4	5
Rational Expectations $X_t = [\mathbb{E}_t[r_{t,t+1}], \mathbb{E}_t[e_{t,t+1}], \mathbb{E}_t[pe_{t,t+1}], \log q_t]'$					
Discount Rate	0.19	0.37	0.57	0.69	0.74
(-) Cash Flow	0.07	0.12	0.19	0.23	0.25
(-) Price-Earnings	0.74	0.51	0.24	0.07	0.02
Residual	0.00	0.00	0.00	0.00	0.00
Subjective Expectations $X_t = [\mathbb{F}_t[r_{t,t+1}], \mathbb{F}_t[e_{t,t+1}], \mathbb{F}_t[pe_{t,t+1}], \log q_t]'$					
Discount Rate	0.02	0.03	0.05	0.05	0.05
(-) Cash Flow	0.34	0.58	0.79	0.90	0.94
(-) Price-Earnings	0.64	0.39	0.16	0.05	0.01
Residual	0.00	0.00	0.00	0.00	0.00

Notes: Table reports variance decompositions based on a Vector Autoregression (VAR) model of the job filling rate and its key determinants under both subjective and rational expectations. The estimated system follows a VAR(1) specification $X_{t+1} = AX_t + \varepsilon_{t+1}$, where the state vector X_t consists of the present discounted values of discount rates, cash flows, log price-earnings ratios, and the log job filling rate. Subjective expectations \mathbb{F}_t are based on survey forecasts of CFOs and IBES financial analysts. Rational expectations \mathbb{E}_t are based on machine learning forecasts from Long Short-Term Memory (LSTM) neural networks. The sample is quarterly from 2005Q1 to 2021Q4.

A.3 Variance Decomposition: First Differences

Table A.3: Variance Decomposition of Aggregate Job Filling Rates: First Differences

Horizon h (Years)	1	2	3	4	5
(a) Rational Expectations: First Differences $\Delta \log q_t = \Delta \mathbb{E}_t[r_{t,t+h}] - \Delta \mathbb{E}_t[e_{t,t+h}] - \Delta \mathbb{E}_t[pe_{t,t+h}]$					
Discount Rate	0.343*	0.479**	0.521**	0.551***	0.587***
t -stat	(1.877)	(2.251)	(2.493)	(2.856)	(3.013)
(-) Cash Flow	0.001	0.027	0.108	0.138	0.151
t -stat	(0.013)	(0.230)	(0.954)	(1.344)	(1.337)
(-) Price-Earnings	0.708**	0.554**	0.448*	0.239	0.218
t -stat	(2.398)	(2.165)	(1.918)	(0.959)	(0.893)
Residual	-0.051	-0.059	-0.077	0.072	0.045
t -stat	(-0.146)	(-0.168)	(-0.232)	(0.217)	(0.136)
N	68	68	68	68	68
(b) Subjective Expectations: First Differences $\Delta \log q_t = \Delta \mathbb{F}_t[r_{t,t+h}] - \Delta \mathbb{F}_t[e_{t,t+h}] - \Delta \mathbb{F}_t[pe_{t,t+h}]$					
Discount Rate	0.028	0.033	0.045	0.074	0.054
t -stat	(1.068)	(0.592)	(0.620)	(0.833)	(0.568)
(-) Cash Flow	0.286*	0.420**	0.601***	0.746***	0.835***
t -stat	(1.779)	(2.517)	(3.443)	(3.884)	(4.346)
(-) Price-Earnings	0.710***	0.529***	0.371	0.204	0.133
t -stat	(2.789)	(2.680)	(1.640)	(0.872)	(0.536)
Residual	-0.023	0.019	-0.016	-0.024	-0.023
t -stat	(-0.076)	(0.071)	(-0.054)	(-0.077)	(-0.070)
N	68	68	68	68	68

Notes: Table reports variance decompositions of the aggregate job filling rate under subjective expectations (panel (a)) or rational expectations (panel (b)). Each row reports the share of the variation in job filling rates that can be explained by h -year present discounted values of discount rates $r_{t,t+h}$, (negative) cash flows $e_{t,t+h}$, and (negative) price-earnings ratios $pe_{t,t+h}$. Residual represents the variation in job filling rates that are not captured by the other components. Subjective expectations \mathbb{F}_t are based on survey forecasts of CFOs and IBES financial analysts. Rational expectations \mathbb{E}_t are based on machine learning forecasts from Long Short-Term Memory (LSTM) neural networks. The sample is quarterly from 2005Q1 to 2021Q4. Newey-West corrected t -statistics with lags = 4 are reported in parentheses: *sig. at 10%. **sig. at 5%. ***sig. at 1%.

A.4 Subjective Decomposition: Additional Controls

The large contribution from subjective long-term cash flow expectations $CF_{\mathbb{F},h}$ in explaining the job filling rate is robust to conditioning on additional variables that could distort the relationship between long-term cash flow expectations and job filling rates. Table A.4 re-estimates the subjective variance decomposition at the 5 year horizon with additional control variables on the right-hand side of the regression: 1 year lag of the log job filling rate, 1 year lag of the dependent variable, and the 1 year ahead survey forecast of the same variable. Controlling for the short-term expectation $\mathbb{F}_t[y_{t+1}]$ accounts for the possibility that survey respondents' long-term forecasts could be influenced by the short-term component of cash flows (Nagel and Xu, 2021).

Table A.4: Variance Decomposition under Subjective Expectations: Additional Controls

Regression: $\mathbb{F}_t[y_{t,t+h}] = \beta_{0,\mathbb{F}} + \beta_{1,\mathbb{F}} \log q_t + \beta_{2,\mathbb{F}} \log q_{t-1} + \beta_{3,\mathbb{F}} \mathbb{F}_{t-1}[y_{t+h-1}] + \beta_{4,\mathbb{F}} \mathbb{F}_t[y_{t+1}] + \varepsilon_{t,\mathbb{F}}$			
Dep Var $y_{t,t+h}$	Discount Rate	Cash Flow	Price-Earnings
Survey \mathbb{F}_t	CFO/IBES	CFO/IBES	CFO/IBES
Horizon h (Years)	5	5	5
β_0	0.035**	0.587***	0.463**
t -stat	(2.495)	(3.051)	(2.329)
β_1	-0.007	-0.734***	0.057
t -stat	(-0.108)	(-3.942)	(0.556)
Adj. R^2	0.456	0.550	0.462
N	68	64	64

Notes: Table reports variance decompositions of the job filling rate under subjective expectations \mathbb{F}_t implied by survey forecasts. $y_{t,t+h}$ denotes the dependent variable of type j to be predicted $h = 5$ years ahead of time t : h year present discounted values of discount rates ($r_{t,t+h} = \sum_{j=1}^h \rho^{j-1} r_{t+j}$), cash flows ($e_{t,t+h} = el_{t+1} + \sum_{j=2}^h \rho^{j-1} \Delta e_{t+j}$), and log price-earnings ratios ($pe_{t,t+h} = \rho^h pe_{t+h}$). Subjective expectations \mathbb{F}_t are based on survey forecasts from the CFO survey for stock returns, and IBES for earnings growth. The sample is quarterly over 2005Q1 to 2019Q3. Newey-West corrected t -statistics with lags = 4 are reported in parentheses: *sig. at 10%. **sig. at 5%. ***sig. at 1%.

A.5 Subjective Cash Flow Expectations from Bloomberg Survey

The large role played by subjective cash flow expectations $CF_{\mathbb{F},h}$ in explaining the job filling rate holds more generally across alternative survey forecasts of earnings growth. Table A.5 re-estimates the subjective variance decomposition while replacing IBES survey forecasts of earnings growth with the corresponding forecast from the Bloomberg (BBG) survey.

Table A.5: Bloomberg (BBG) Survey Earnings Expectations

Horizon h (Years)	1	2	3	4	5
Subjective Expectations $\mathbb{F}_t[\cdot]$ $\log q_t = c_q + \mathbb{F}_t[r_{t,t+h}] - \mathbb{F}_t[e_{t,t+h}] - \mathbb{F}_t[pe_{t,t+h}]$					
Discount Rate	0.012	-0.005	0.010	0.037	0.009
t -stat	(0.974)	(-0.130)	(0.192)	(0.566)	(0.180)
(-) Cash Flow	0.586***	0.830***	0.851***	0.896***	0.949***
t -stat	(8.476)	(8.317)	(7.213)	(5.288)	(4.541)
(-) Price-Earnings	0.389***	0.202***	0.154***	0.093***	0.081**
t -stat	(6.008)	(6.386)	(4.712)	(3.998)	(2.422)
Residual	0.028	0.024	0.015	0.039	0.014
t -stat	(0.440)	(0.415)	(0.178)	(0.337)	(0.063)

Notes: Table reports variance decompositions of the job filling rate while replacing IBES earnings growth forecast with BBG survey as measures of subjective cash flows. The sample is quarterly from 2006Q1 to 2022Q4. Newey-West corrected t -statistics with lags = 4 are reported in parentheses: *sig. at 10%. **sig. at 5%. ***sig. at 1%.

A.6 Alternative Survey Expectations of Cash Flows: CFO Survey

The large role played by subjective cash flow expectations $CF_{\mathbb{F},h}$ in explaining the job filling rate holds more generally across alternative survey forecasts of earnings growth. Table A.6 re-estimates the subjective variance decomposition while replacing 1 year ahead IBES survey forecasts of earnings growth with the corresponding forecast from the CFO survey. The estimates suggest a significant negative relationship between subjective cash expectations and the job filling rate.

Subjective expectations of log price-earnings $\mathbb{F}_t[pe_{t+h}]$ combine the current log price-earnings ratio pe_t with subjective expectations of the stock returns $\mathbb{F}_t[r_{t+h}]$ and earnings growth $\mathbb{F}_t[\Delta e_{t+h}]$, both from the same set of respondents to the CFO survey.

$$\mathbb{F}_t[pe_{t+h}] = \frac{1}{\rho^h} pe_t - \frac{1}{\rho^h} \sum_{j=1}^h \rho^{j-1} (c_{pe} + \underbrace{\mathbb{F}_t[\Delta e_{t+j}]}_{\text{Survey (CFO)}} - \underbrace{\mathbb{F}_t[r_{t+j}]}_{\text{Survey (CFO)}}) \quad (\text{A.2})$$

Subjective expectations of log earnings-employment $\mathbb{F}_t[el_{t+1}]$ combine the current log earnings-employment ratio el_t with subjective CFO expectations of earnings growth $\mathbb{F}_t[\Delta e_{t,t+1}]$:

$$\mathbb{F}_t[el_{t+1}] = el_t + \underbrace{\mathbb{F}_t[\Delta e_{t,t+1}]}_{\text{Survey (CFO)}} \quad (\text{A.3})$$

The forecast horizon has been limited to $h = 1$ year ahead and the sample covers a shorter period over 2005Q1 to 2019Q3 due to missing earnings growth forecasts in the CFO survey.

Table A.6: Variance Decomposition under Subjective Expectations: CFO Survey

Regression: $\mathbb{F}_t[y_{t,t+h}] = \beta_{0,\mathbb{F}} + \beta_{1,\mathbb{F}} \log q_t + \varepsilon_{t,\mathbb{F}}$			
Dep Var $y_{t,t+h}$	Discount Rate	Cash Flow	Price-Earnings
Survey \mathbb{F}_t	CFO	CFO	CFO
Horizon h (Years)	1	1	1
β_0	0.034***	1.302***	2.117***
t -stat	(9.917)	(10.000)	(20.583)
β_1	0.007	-0.637*	-0.373**
t -stat	(0.325)	(-1.934)	(-2.467)
Adj. R^2	0.002	0.081	0.097
N	59	59	59

Notes: Table reports variance decompositions of the job filling rate under subjective expectations \mathbb{F}_t implied by survey forecasts from the CFO survey. $y_{t,t+h}$ denotes the dependent variable of type j to be predicted $h = 1$ years ahead of time t : h year present discounted values of discount rates ($r_{t,t+h} = \sum_{j=1}^h \rho^{j-1} r_{t+j}$), cash flows ($e_{t,t+h} = el_{t+1} + \sum_{j=2}^h \rho^{j-1} \Delta e_{t+j}$), and log price-earnings ratios ($pe_{t,t+h} = \rho^h pe_{t+h}$). The sample is quarterly over 2005Q1 to 2019Q3. Newey-West corrected t -statistics with lags = 4 are reported in parentheses: *sig. at 10%. **sig. at 5%. ***sig. at 1%.

A.7 Alternative Survey Expectations of Stock Returns

The small role played by subjective discount rate expectations in explaining the job filling rate holds more generally across alternative survey forecasts of stock returns. Table A.7 reports estimates from regressing 1 year ahead survey expectations of stock returns $\mathbb{F}_t[r_{t,t+h}]$ on the log job filling rate q_t under alternative survey forecasts of stock returns. In all survey measures, the estimates suggest a weak relationship between subjective stock return expectations $\mathbb{F}_t^s[r_{t,t+h}]$ and the job filling rate q_t .

$r_{t,t+h}$ denotes h year CRSP stock returns (with dividends) or S&P 500 price growth from time t to $t + h$, depending on the concept that survey respondents are asked to predict: log stock returns for CB, SOC, Gallup/UBS, and CFO; log price growth for Livingston. $\mathbb{F}_t^s[r_{t,t+h}]$ denotes subjective expectations of stock returns or price growth from survey s . CoC and Hurdle denotes corporate cost of capital and hurdle rates constructed in Gormsen and Huber (2023). The forecast horizon has been limited to 1 year ahead due to limited data availability in the alternative surveys. The sample is quarterly over 2005Q1 to 2021Q4 when considering the NX, CB, SOC, and CFO surveys, 2005Q1 to 2007Q4 for Gallup/UBS, and semi-annual over 2005Q1 to 2021Q4 from Q2 and Q4 of each calendar year for Livingston.

Table A.7: The Role of Subjective Discount Rates: Alternative Surveys

Dep Var	Regression: $\mathbb{F}_t^s[r_{t,t+h}] = \beta_{0,\mathbb{F}} + \beta_{1,\mathbb{F}} \log q_t + \varepsilon_{t,\mathbb{F}}$							
	Discount Rate	Discount Rate	Discount Rate	Discount Rate	Discount Rate	Discount Rate	Discount Rate	Discount Rate
Survey s	CFO	NX	CB	SOC	Gallup	Liv	CoC	Hurdle
Horizon h (Years)	1	1	1	1	1	1	1	1
β_0	0.034***	0.061***	0.041***	0.042***	0.062***	0.048***	0.060***	0.083***
t -stat	(20.053)	(38.202)	(33.273)	(43.401)	(30.778)	(6.397)	(58.887)	(145.982)
β_1	0.013	-0.011	0.026**	0.002	-0.065***	0.067	0.024***	0.013**
t -stat	(1.039)	(-0.995)	(2.015)	(0.414)	(-3.688)	(0.722)	(2.936)	(2.087)
Adj. R^2	0.029	0.012	0.069	0.009	0.216	0.045	0.232	0.154
N	68	68	68	68	16	34	68	68

Notes: Table reports intercept (β_0) and slope (β_1) estimates from regressing $h = 1$ year ahead survey expectations of stock returns $\mathbb{F}_t[r_{t,t+h}]$ on the log job filling rate q_t . $r_{t,t+h}$ denotes h year CRSP stock returns (with dividends) or S&P 500 price growth from time t to $t + h$, depending on the concept that survey respondents are asked to predict: log stock returns for CB, SOC, Gallup/UBS, and CFO; log price growth for Livingston. $\mathbb{F}_t^s[r_{t,t+h}]$ denotes subjective expectations of stock returns or price growth from survey s . CoC and Hurdle denotes corporate cost of capital and hurdle rates constructed in Gormsen and Huber (2023). The sample is quarterly over 2005Q1 to 2021Q4 when considering the NX, CB, SOC, and CFO surveys, 2005Q1 to 2007Q4 for Gallup/UBS, and semi-annual over 2005Q1 to 2021Q4 from Q2 and Q4 of each calendar year for Livingston. Newey-West corrected t -statistics with lags = 4 are reported in parentheses: *sig. at 10%. **sig. at 5%. ***sig. at 1%.

A.8 Bartik Instrument for Subjective Beliefs

A key challenge in estimating the regional decomposition is that regional labor market tightness and subjective beliefs are endogenously determined jointly, which likely overstates the estimated effect of subjective beliefs on cash flows. Another issue is the limited availability of subjective expectations data at a regional level, as firms in the data could operate across multiple regions.

I address these challenges by constructing a Bartik shift-share instrument $\widetilde{\mathbb{F}}_{s,t}[\cdot]$ to proxy for regional subjective expectations $\mathbb{F}_t[\cdot]$:

$$\widetilde{\mathbb{F}}_{s,t}[y_{s,t,t+h}] = \sum_{i \in I} s_{s,i,t} \mathbb{F}_t[y_{i,t,t+h}], \quad y = r, e, pe \quad (\text{A.4})$$

where time t is measured quarterly, and $s_{s,i,t}$ is the employment share of industry i in region s , sourced from the BLS Quarterly Census of Employment and Wages (QCEW). I measure industries as 2-digit NAICS sectors and regions as U.S. states.

$\mathbb{F}_t[y_{i,t,t+h}]$ denotes aggregate subjective expectations for variable y in industry i , where the variables are either discount rates, cash flows, or price-earnings ratios. I construct these industry-level expectations by aggregating the firm-level expectations across all firms in industry i from the IBES database. I construct earnings growth using earnings per share (EPS) forecasts, discount rates from price target forecasts, and employment growth by projecting the firm's past realized employment growth onto the subjective forecast from the CFO survey. I measure regional labor market tightness at the state level as the vacancy-to-unemployment ratio, where vacancies use data from the Job Openings and Labor Turnover Survey (JOLTS) and unemployment uses data from the Local Area Unemployment Statistics (LAUS).

Using this Bartik instrument, I estimate the following regression:

$$\theta_{s,t} = \beta \cdot \widetilde{\mathbb{F}}_t[y_{s,t,t+1}] + \alpha_s + \alpha_t + \varepsilon_t, \quad y = r, e, pe \quad (\text{A.5})$$

where $\widetilde{\mathbb{F}}_t[y_{s,t,t+1}]$ is the Bartik shift-share instrument for subjective expectations from equation (A.4), α_i are region fixed effects, and α_t are time fixed effects. The region fixed effects α_s can absorb regional differences in vacancy posting cost $c_{\theta,s}$, and time fixed effects α_t can absorb long-run inflation expectations common across regions within a common monetary union (Hazell et al., 2022). The identification of the parameter β hinges on the quasi-random assignment of industry-level shocks, implying that these shocks are, in expectation, uncorrelated with relevant unobservables. However, identification could be threatened if preexisting trends in labor market tightness were more prevalent in states with industry mixes that would make them more or less susceptible to the shock.

The estimates in panel (a) columns (4)-(6) of Table A.8 suggest that regional labor market tightness is primarily driven by subjective cash flow news rather than discount rate news. Discount rate expectations play a smaller but still sizeable role. At the one-year horizon, subjective cash flow news accounts for 58.8% of the variation, while discount rate news contributes 2.7%. These estimates are in line with the aggregate estimates from Table 2. The sign of the

relationship to labor market tightness is also consistent with predictions from the search model. Based on the regional model, higher labor market tightness predicts lower discount rates, higher cash flows, or both. Panel (b) columns (4)-(6) condition on time and region fixed effects to find that cash flow news plays an even stronger role in explaining regional labor market fluctuations. At the one-year horizon, subjective cash flow news accounts for 97.8% of the variation, while discount rate news contributes 1.1%.

Table A.8: Determinants of Regional Labor Market Tightness: Bartik Instrument for Subjective Beliefs

	$\theta_{s,t} = \beta \cdot \tilde{\mathbb{F}}_t[y_{s,t,t+1}] + \alpha_s + \alpha_t + \varepsilon_t, y = r, e, pe$					
	(1) Discount Rate	(2) Cash Flow	(3) Price Earning	(4) Discount Rate	(5) Cash Flow	(6) Price Earning
$\beta^{\mathbb{F},y}$	-0.0705** (0.0301)	0.4496*** (0.0299)	0.5153*** (0.0341)	0.0270 (0.0299)	0.5884*** (0.0301)	0.4578*** (0.0130)
Observations	12,291	12,291	12,291	12,291	12,291	12,291
R^2	0.98	0.24	0.19	0.97	0.30	0.22
Specification	OLS	OLS	OLS	Bartik	Bartik	Bartik
Region FE	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	No	No	No	No	No	No

Notes: Subjective expectations \mathbb{F}_t are based on survey forecasts of CFOs and IBES financial analysts. Rational expectations \mathbb{E}_t are based on machine learning forecasts from Long Short-Term Memory (LSTM) neural networks. The sample is quarterly from 2005Q1 to 2022Q4. Observations weighted by size of labor force in each state. Standard errors clustered by state are reported in parentheses: *sig. at 10%. **sig. at 5%. ***sig. at 1%.

A.9 Price-Employment Ratios and Job Filling Rates

The job filling rate in the search model is closely related to price-employment ratio. Under constant returns to scale, the hiring equation directly relates the job filling rate with the price-employment ratio:

$$\frac{Cov [pl_t, \log q_t]}{Var [\log q_t]} \approx -1$$

where pl_t is the log stock price to employment ratio. Table A.9 tests this prediction by regressing log price-employment on log job filling rate:

$$pl_t = \alpha + \beta \log q_t + \varepsilon_t$$

The estimates show that β not significantly different from -1, and the R^2 of the regression is relatively high, consistent with predictions of the model.

Table A.9: Role of Price-Employment Ratios

Dep. Var. Specification	Log Price-Employment pl_t	
	Levels	First Differences
Share of job filling rate variation (β)	-0.988	-0.977
t -stat	(0.204)	(0.203)
Adj. R^2	0.895	0.780
N	76	76

Notes: Table reports slope coefficient from regressing the aggregate log price-employment ratio on the aggregate log job filling rate. The sample is quarterly from 2002Q1 to 2020Q4. Newey-West corrected t -statistics for a two-sided t -test of $H_0 : \beta = -1$ with lags = 4 are reported in parentheses: *sig. at 10%. **sig. at 5%. ***sig. at 1%.

A.10 Price-Earnings Ratios and Job Filling Rates

To examine the role of price-earnings ratios (PE ratios) directly, I decompose the job filling rate using equation (6), before applying the Campbell and Shiller (1988) decomposition of the price-earnings ratio pe_t :

$$\log q_t = \log \kappa - pe_t + \Delta e_{t+1} - el_{t+1}$$

Table A.10 reports the estimates. Under rational expectations, PE ratios explain nearly all the variation in job filling rates (101.9%), earnings growth is positively related to the job filling rate. Under subjective expectations, the role of earnings growth becomes negative, suggesting that survey respondents systematically over-react to news about cash flows. These findings confirm that observed hiring decisions reflect and over-weighting of cash flow expectations, distorting the relationship between labor market fluctuations and firm fundamentals.

Table A.10: Role of Price-Earnings Ratios

Dep. Var. Horizon h (Years)	Price- Earnings 1	Earnings Growth 1	Earnings- Employment 1	Residual 1
Rational Expectations $\log q_t = \log \kappa - pe_t + \mathbb{E}_t[\Delta e_{t+1}] - \mathbb{E}_t[el_{t+1}]$				
Share of job filling rate variation t -stat	-1.019*** (-3.569)	0.250*** (8.189)	-0.253*** (-8.141)	-0.022 (-0.221)
Subjective Expectations $\log q_t = \log \kappa - pe_t + \mathbb{F}_t[\Delta e_{t+1}] - \mathbb{F}_t[el_{t+1}]$				
Share of job filling rate variation t -stat	-1.019*** (-3.569)	0.644*** (8.842)	-0.586*** (-8.476)	0.039 (0.418)

Notes: Table decomposes the job filling rate using equation (6), before applying the Campbell and Shiller (1988) decomposition of the price-earnings ratio pe_t . Subjective expectations \mathbb{F}_t are based on survey forecasts of CFOs and IBES financial analysts. Rational expectations \mathbb{E}_t are based on machine learning forecasts from Long Short-Term Memory (LSTM) neural networks. The sample is quarterly from 2005Q1 to 2021Q4. Newey-West corrected t -statistics with lags = 4 are reported in parentheses: *sig. at 10%. **sig. at 5%. ***sig. at 1%.

A.11 Variance Decomposition of Aggregate Labor Market Tightness

Table A.11 estimates a variance decomposition of aggregate labor market tightness (θ_t), defined as the ratio of vacancies to unemployment, under subjective expectations and rational expectations. Under rational expectations, discount rates and cash flows both contribute to fluctuations in labor market tightness. Under subjective expectations, cash flow news dominates, explaining nearly 80% of the variation in labor market tightness. Fluctuations in subjective discount rates play no significant role. These results suggest that belief distortions in cash flow expectations are a primary driver of aggregate labor market dynamics, leading to excessive volatility in unemployment and hiring decisions.

Table A.11: Variance Decomposition of Labor Market Tightness

Horizon h (Years)	1	2	3	4	5
Rational Expectations					
$\theta_t = c_\theta - \mathbb{E}_t[r_{t,t+h}] + \mathbb{E}_t[e_{t,t+h}] + \mathbb{E}_t[pe_{t,t+h}]$					
(-) Discount Rate	0.131***	0.275***	0.435***	0.475***	0.515***
t -stat	(3.478)	(3.577)	(3.793)	(3.045)	(2.678)
Cash Flow	0.060	0.133	0.149	0.159	0.171*
t -stat	(1.331)	(1.560)	(1.376)	(0.836)	(1.722)
Price-Earnings	0.795***	0.617***	0.450	0.371	0.288
t -stat	(4.452)	(2.751)	(1.168)	(0.935)	(0.555)
Residual	0.013	-0.025	-0.033	-0.006	0.026
t -stat	(0.069)	(-0.099)	(-0.080)	(-0.013)	(0.045)
N	68	68	68	68	68
Subjective Expectations					
$\theta_t = c_\theta - \mathbb{F}_t[r_{t,t+h}] + \mathbb{F}_t[e_{t,t+h}] + \mathbb{F}_t[pe_{t,t+h}]$					
(-) Discount Rate	0.015	0.018	0.033	-0.006	-0.094
t -stat	(1.128)	(0.771)	(0.940)	(-0.090)	(-1.060)
Cash Flow	0.688***	0.698***	0.666***	0.737***	0.772***
t -stat	(5.599)	(5.772)	(4.586)	(4.182)	(4.623)
Price-Earnings	0.331**	0.253	0.287	0.237	0.322
t -stat	(2.044)	(1.420)	(0.882)	(0.726)	(1.124)
Residual	-0.034	0.032	0.014	0.032	0.001
t -stat	(-0.165)	(0.146)	(0.038)	(0.086)	(0.002)
N	68	68	68	68	68

Notes: Table reports variance decompositions of aggregate labor market tightness, defined as the ratio of vacancies to unemployment, under subjective expectations or rational expectations. Subjective expectations \mathbb{F}_t are based on survey forecasts of CFOs and IBES financial analysts. Rational expectations \mathbb{E}_t are based on machine learning forecasts from Long Short-Term Memory (LSTM) neural networks. The sample is quarterly from 2005Q1 to 2021Q4. Newey-West corrected t -statistics with lags = 4 are reported in parentheses: *sig. at 10%. **sig. at 5%. ***sig. at 1%.

A.12 Biases in Subjective Beliefs

To directly quantify the importance of biases in subjective beliefs, I consider predictive regressions of biases in subjective expectations of discount rates, cash flows, and price-earnings ratios on the job filling rate. I define the bias as the difference between subjective and machine expectations. Table A.12 reports estimates $\beta_{1,B}$ from regressing biases in subjective discount rate, cash flow, and log price-earnings expectations on the job filling rate:

$$Bias_t[y_{t,t+h}] = \beta_{0,B} + \beta_{1,B} \log q_t + \varepsilon_{t,B}, \quad y = r, e, pe \quad (\text{A.6})$$

where the $Bias_t[y_{t,t+h}] \equiv \mathbb{F}_t[y_{t,t+h}] - \mathbb{E}_t[y_{t,t+h}]$ is defined as the difference between subjective and machine expectations of the same variable.

The results indicate that biases in survey forecasts are important contributors to fluctuations in job filling rates, especially at longer horizons. At the 5-year horizon, biases in cash flow expectations lead survey respondents to over-weight 70.5% of the variation in job filling rates to the cash flow component. This mis-perception is counteracted by biases in subjective discount rate expectations, which leads survey respondents to under-weight 67.5% of the variation in the job filling rate. These findings emphasize the importance of belief distortions in driving labor market fluctuations. The profile of the response across forecast horizons is broadly consistent with the profile of the MSE ratios across horizons in Table 1. For discount rate and cash flow expectations, the machine outperformed the survey by a wider margin over longer horizons, suggesting that the bias in survey responses likely play a bigger role over these longer horizons.

Table A.12: Biases in Subjective Beliefs and the Job Filling Rate

Horizon h (Years)	1	2	3	4	5
Biases in Subjective Expectations					
	$\mathbb{F}_t[y_{t,t+h}] - \mathbb{E}_t[y_{t,t+h}] = \beta_{0,B} + \beta_{1,B} \log q_t + \varepsilon_{t,B}, \quad y = r, e, pe$				
Discount Rate	-0.169	-0.291**	-0.543**	-0.603***	-0.675***
t -stat	(-1.474)	(-1.992)	(-2.505)	(-2.734)	(-2.876)
(-) Cash Flow	0.291	0.543***	0.615***	0.649***	0.705***
t -stat	(1.162)	(3.064)	(4.053)	(4.025)	(3.553)
(-) Price-Earnings	-0.084	-0.146	-0.043	-0.127	-0.094
t -stat	(-0.246)	(-1.311)	(-0.195)	(-0.439)	(-0.295)
Residual	0.037	0.106	0.029	-0.081	-0.063
t -stat	(0.085)	(0.416)	(0.085)	(-0.204)	(-0.143)
N	68	68	68	68	68

Notes: This table reports estimates $\beta_{1,B}$ from regressing the survey bias $\mathbb{F}_t[y_{t,t+h}] - \mathbb{E}_t[y_{t,t+h}]$ on the job filling rate q_t . $y_{t,t+h}$ denotes the dependent variable of type j to be predicted h years ahead of time t . The components of the decomposition are h -year present discounted values of discount rates $r_{t,t+h}$, (negative) cash flows $e_{t,t+h}$, and (negative) price-earnings ratios $pe_{t,t+h}$. The residual term captures variation in the bias that cannot be explained by the three components. Subjective expectations \mathbb{F}_t are based on survey forecasts from the CFO survey for stock returns, and IBES for earnings growth. Machine expectations are based on machine learning forecasts \mathbb{E}_t from Long Short-Term Memory (LSTM) neural networks $G(\mathcal{X}_t, \beta_{h,t})$, whose parameters $\beta_{h,t}$ are estimated in real time using \mathcal{X}_t , a large scale dataset of macroeconomic, financial, and textual data. The bias is defined as the difference between subjective and machine expectations: $Bias_t = \mathbb{F}_t - \mathbb{E}_t$. The sample is quarterly from 2005Q1 to 2021Q4. Newey-West corrected t -statistics with lags = 4 are reported in parentheses: *sig. at 10%. **sig. at 5%. ***sig. at 1%.

A.13 Risk-Neutral Measure

Table A.13 estimates the variance decomposition of job filling rates under risk-neutral probabilities, defined as:

$$\mathbb{E}_t[M_{t+1}X_{t+1}] = \frac{1}{R_{f,t}}\mathbb{E}_t^Q[X_{t+1}] \quad (\text{A.7})$$

$$\mathbb{F}_t[M_{t+1}X_{t+1}] = \frac{1}{R_{f,t}}\mathbb{F}_t^Q[X_{t+1}] \quad (\text{A.8})$$

where $\mathbb{E}_t^Q[\cdot]$ and $\mathbb{F}_t^Q[\cdot]$ denote rational and subjective expectations computed under risk-neutral probabilities. I measure risk-free rates $R_{f,t}$ using the 3-month Treasury bill rate.

Under rational expectations, risk-neutral discount rate fluctuations explain an insignificant share (0.1%) of the job filling rate, while the terminal value from future price-earnings ratios account for 80.7%. Under subjective expectations, the weight on cash flows increases even further (97.7%) with virtually no role for discount rates. These results reinforce the idea that cash flow over-reaction is a key feature of subjective expectations and lead to excessive volatility in job filling rates.

Table A.13: Risk-Neutral Expectations

Dep. Var. Horizon h (Years)	Discount Rate 5	(-) Cash Flow 5	(-) Price-Earnings 5	Residual 5
Risk-Neutral Expectations (Rational) $\mathbb{E}_t[M_{t+1}X_{t+1}] = \frac{1}{R_{f,t}}\mathbb{E}_t^Q[X_{t+1}]$ $\log q_t = c_q + \mathbb{E}_t^Q[r_{t,t+h}] - \mathbb{E}_t^Q[e_{t,t+h}] - \mathbb{E}_t^Q[pe_{t,t+h}]$				
Share of $Var[\log q_t]$ t -stat	0.001 (0.040)	0.203*** (2.834)	0.807** (2.493)	-0.010
Risk-Neutral Expectations (Subjective) $\mathbb{F}_t[M_{t+1}X_{t+1}] = \frac{1}{R_{f,t}}\mathbb{F}_t^Q[X_{t+1}]$ $\log q_t = c_q + \mathbb{F}_t^Q[r_{t,t+h}] - \mathbb{F}_t^Q[e_{t,t+h}] - \mathbb{F}_t^Q[pe_{t,t+h}]$				
Share of $Var[\log q_t]$ t -stat	-0.004 (-0.378)	0.977*** (5.075)	-0.023*** (3.596)	0.050

Notes: Subjective expectations \mathbb{F}_t are based on survey forecasts of CFOs and IBES financial analysts. Rational expectations \mathbb{E}_t are based on machine learning forecasts from Long Short-Term Memory (LSTM) neural networks. Risk-free rates are measured $R_{f,t}$ using the 3-month Treasury bill rate. The sample is quarterly from 2016Q1 to 2022Q4. Newey-West corrected t -statistics with lags = 4 are reported in parentheses: *sig. at 10%. **sig. at 5%. ***sig. at 1%.

A.14 Subjective Risk-Neutral Measure from Dividend Futures

To validate the robustness of the risk-neutral results, I construct risk-neutral cash flow expectations by using dividend futures following Gormsen and Kojen (2020).

The estimates in Table A.14 show that subjective cash flow news remains the dominant driver of the job filling rate. Compared to survey-based subjective expectations, the subjective risk-neutral expectation places more weight on long-horizon cash flows, since the price-earnings ratio component (terminal value) is larger than in Table A.13. Risk-free rates play a minimal role, suggesting that fluctuations in job filling rates are more closely tied to belief distortions about earnings growth rather than movements in risk premia. These findings support the view that managerial expectations about future earnings are systematically distorted, with significant over-reaction to long-term cash flow news.

Table A.14: Subjective Risk-Neutral Expectations from Dividend Futures

Dep. Var. Horizon h (Years)	Discount Rate 1	(-) Cash Flow 1	(-) Price-Earnings 1	Residual 1
Risk-Neutral Expectations $\mathbb{F}_t^Q[\cdot]$ $\log q_t = c_q + \mathbb{F}_t^Q[r_{t,t+h}] - \mathbb{F}_t^Q[e_{t,t+h}] - \mathbb{F}_t^Q[pe_{t,t+h}]$				
Share of $Var[\log q_t]$	-0.003	0.430***	0.553***	0.019
t -stat	(-1.583)	(10.338)	(4.704)	(0.101)

Notes: Risk-neutral expectations \mathbb{F}_t^Q about discount rates are based on the 3-month Treasury bill rate. Risk-neutral expectations \mathbb{F}_t^Q about cash flows are based on dividend futures following Gormsen and Kojen (2020), which are fitted values from regressing realized 1-year dividend growth $\Delta d_{h,t}$ on the S&P 500 on the 1-year equity yield $e_{h,t}$ of the same index: $\Delta d_{h,t} = \beta_0 + \beta_1 e_{h,t} + \epsilon_t$ where t is measured in quarters, h is the horizon of the expectation, the coefficients are estimated recursively using expanding samples up to the time of each observation. The equity yield is defined as $e_{h,t} = (1/h) \log(D_t/F_{h,t})$, where D_t is the quarterly dividend on the S&P 500 and $F_{h,t}$ is the price of 1-year ahead dividend futures from Bloomberg. The sample is quarterly from 2016Q1 to 2022Q4. Newey-West corrected t -statistics with lags = 4 are reported in parentheses: *sig. at 10%. **sig. at 5%. ***sig. at 1%.

A.15 Ex-post Decomposition

Since the log-linear decomposition of the job filling rate holds both ex-ante and ex-post, a variance decomposition of the job filling rate can also be estimated using ex-post realized data, under the assumption of the manager's perfect foresight:

$$1 \approx \underbrace{\frac{Cov[r_{t,t+h}, \log q_t]}{Var[\log q_t]}}_{\text{Discount Rate news}} - \underbrace{\frac{Cov[e_{t,t+h}, \log q_t]}{Var[\log q_t]}}_{\text{Cash Flow News}} - \underbrace{\frac{Cov[pe_{t,t+h}, \log q_t]}{Var[\log q_t]}}_{\text{Future Price-Earnings News}}$$

Table A.15 reports the estimates. Ex-post discount rate news is the main driver of fluctuations in job filling rates, with ex-post cash flow news playing only a minor role. At the 5 year horizon, 79.4% of the variation in the job filling rate is driven by discount rate news. In contrast, cash flow news has a smaller effect, contributing only 10.3% over the same period. The ex-post decomposition differs from both rational and subjective decompositions, suggesting that survey respondents not only need to learn about the true data generating process of the economy (when compared against rational expectations), but also hold significant distortions in their beliefs about them (when compared against subjective expectations).

Table A.15: Ex-Post Variance Decomposition of the Job Filling Rate

Horizon h (Years)	1	2	3	4	5
Ex-Post Decomposition (Full Information Rational Expectations)					
$\log q_t = c_q + r_{t,t+h} - e_{t,t+h} - pe_{t,t+h}$					
Discount Rate	0.445	0.675**	0.705***	0.780***	0.794***
t -stat	(1.220)	(2.429)	(2.816)	(3.850)	(3.439)
(-) Cash Flow	-0.017	0.047	0.083	0.080	0.103
t -stat	(-0.191)	(0.567)	(0.996)	(1.026)	(1.206)
(-) Price-Earnings	0.572***	0.273***	0.208***	0.123	0.074
t -stat	(7.423)	(5.965)	(4.037)	(1.468)	(0.671)
Residual	-0.000	0.005	0.003	0.016	0.028
t -stat	(-0.001)	(0.031)	(0.017)	(0.089)	(0.137)
N	68	68	68	68	68

Notes: Table reports variance decompositions of the job filling rate from equation using ex-post realized outcomes. The components of the decomposition are present discounted values of discount rates $r_{t,t+h}$, cash flows $e_{t,t+h}$, and log price-earnings ratios $pe_{t,t+h}$. Residual represents the variation in job filling rates that are not captured by the other components. The sample is quarterly from 2005Q1 to 2021Q4. Newey-West corrected t -statistics with lags = 4 are reported in parentheses: *sig. at 10%. **sig. at 5%. ***sig. at 1%.

A.16 Counter-Cyclical Dispersion in Subjective Expectations

Figure A.1: Counter-Cyclical Dispersion in Subjective Expectations



Notes: This figure plots the cross-sectional dispersion of subjective expectations (Survey) for each outcome that decomposes the job filling rate (left axis), against the current detrended job filling rate q_t (right axis). The dispersion is calculated as the standard deviation of each outcome across firms covered in the IBES database, weighted by the firm's market capitalization. The components of the decomposition are $h = 5$ year present discounted values of discount rates $r_{t,t+h}$, (negative) cash flows $e_{t,t+h}$, and (negative) price-earnings ratios $pe_{t,t+h}$. x axis denotes the date on which actual values and job filling rates were realized, and the vertical distance between the actual line and each forecast is the forecast error. Subjective expectations \mathbb{F}_t are based on survey forecasts from the CFO survey for stock returns, and IBES for earnings growth. Machine expectations are based on machine learning forecasts \mathbb{E}_t from Long Short-Term Memory (LSTM) neural networks $G(\mathcal{X}_t, \beta_{h,t})$, whose parameters $\beta_{h,t}$ are estimated in real time using \mathcal{X}_t , a large scale dataset of macroeconomic, financial, and textual data. The out-of-sample testing period for survey and machine forecasts is quarterly and spans 2005Q1 to 2021Q4. NBER recessions are shown with gray shaded bars.

B Model Details

B.1 Representative Agent Model

In this section, I present a search and matching model based on Diamond (1982), Mortensen (1982), and Pissarides (2009). The model introduces subjective beliefs that may depart from rational expectations, thereby capturing the impact of belief distortions on labor market dynamics. See Petrosky-Nadeau et al. (2018) for a standard search and matching model formulated under rational expectations.

Environment Consider a discrete time economy populated by a representative household and a representative firm that uses labor as a single input to production.

Representative Household The household has a continuum of mass 1 members who are either employed L_t or unemployed U_t at any point in time. The population is normalized to 1, i.e., $L_t + U_t = 1$, meaning that L_t and U_t are also the rates of employment and unemployment, respectively. The household's consumption decision implies a stochastic discount factor M_{t+1} . The household pools the income of all members before making its consumption decision. Assume that the household's members have access to complete contingent claims against aggregate risk. Assume that the household has perfect consumption insurance. Risk sharing implies each member consumes the same amount regardless of idiosyncratic shocks.

Search and Matching We adopt a standard end-of-period matching convention (Petrosky-Nadeau et al., 2018). At the start of period t , the employment stock L_t reflects the total number of workers carried over from the previous period before any separations or new hires in period t . A fraction δ_t of these workers separate during the period, so the number of continuing employees becomes $(1 - \delta_t)L_t$.

The representative firm posts job vacancies V_t and engage in search over the course of the period to attract unemployed workers U_t . Matches are formed at the end of period t according to a matching function $m(U_t, V_t)$, where $q_t \equiv m(U_t, V_t)/V_t$ is the *job filling rate*, and $f_t \equiv m(U_t, V_t)/U_t$ is the *job finding rate*. These new matches become part of the workforce starting in period $t + 1$, so employment evolves according to the employment accumulation equation:

$$L_{t+1} = (1 - \delta_t)L_t + q_t V_t \tag{A.9}$$

The job filling rate q_t maps vacancy posting decisions made during period t into employment outcomes observed at the beginning of period $t + 1$. The variance decomposition does not require specifying the functional form of the matching function m . As explained in Section D, I measure the number of matches directly from the data using changes in unemployment and the number of short-term unemployed (Shimer, 2012). Posting a vacancy costs the firm $\kappa > 0$ per period,

reflecting fixed hiring costs such as training and administrative setup. Jobs are destroyed at a time-varying job separation rate δ_t . Unemployment $U_t = 1 - L_t$ evolves according to:

$$U_{t+1} = \delta_t(1 - U_t) + (1 - q_t\theta_t)U_t \quad (\text{A.10})$$

where $\theta_t = V_t/U_t$ denotes labor market tightness, defined as the vacancy-to-unemployment ratio.

Representative Firm The firm has access to a production function F which uses labor L_t as an input to produce output $Y_t = F(L_t)$. Dividends to the firm's shareholders E_t are defined as $E_t \equiv \Pi_t - \kappa V_t$, where $\Pi_t \equiv Y_t - W_t L_t$ is the total profit before vacancy posting costs κV_t and W_t is the wage rate. As in Petrosky-Nadeau et al. (2018), I assume that the representative household owns the equity of the firm, and that the firm pays out all of its earnings as dividends. I also assume that firms have the same unconstrained access to financing as investors in the financial market. The firm posts the optimal number of vacancies to maximize the cum-dividend market value of equity S_t :

$$S_t = \max_{\{V_{t+j}, L_{t+j}\}_{j=0}^{\infty}} \mathbb{F}_t \left[\sum_{j=0}^{\infty} M_{t,t+j} E_{t+j} \right] \quad (\text{A.11})$$

subject to the employment accumulation equation (??). The firm takes the wage rate W_t , household's stochastic discount factor $M_{t,t+j} = \prod_{s=1}^j M_{t+s}$, and job filling rate q_t as given.

Expectations Let $\mathbb{F}_t[\cdot]$ denote expectations conditional on information available at the beginning of period t , computed based on the manager's possibly distorted beliefs. I allow for departures from rational expectations $\mathbb{E}_t[\cdot]$, but restrict to beliefs that preserve the law of iterated expectations.

Hiring Equation The firm's optimal hiring decision implies the *hiring equation*, which equates the expected discounted value of hiring a marginal worker with its marginal cost. Rewrite the infinite-horizon value-maximization problem of the firm into recursive form:

$$S_t = \max_{V_t, L_{t+1}} \Pi_t - \kappa V_t + \mathbb{F}_t [M_{t+1} S_{t+1}] \quad (\text{A.12})$$

$$\text{s.t. } L_{t+1} = (1 - \delta_t)L_t + q_t V_t \quad (\text{A.13})$$

The first-order condition with respect to V_t is:

$$\frac{\partial S_t}{\partial V_t} = -\kappa + \mathbb{F}_t \left[M_{t+1} \frac{\partial S_{t+1}}{\partial L_{t+1}} \frac{\partial L_{t+1}}{\partial V_t} \right] = 0 \quad (\text{A.14})$$

Substitute $\frac{\partial L_{t+1}}{\partial V_t} = q_t$ and $\frac{\partial L_{t+1}}{\partial L_t} = (1 - \delta_t)$ from the employment accumulation equation (A.13), and rearrange (A.14) in terms of the marginal cost of hiring κ/q_t :

$$\frac{\kappa}{q_t} = \mathbb{F}_t \left[M_{t+1} \frac{\partial S_{t+1}}{\partial L_{t+1}} \right] \quad (\text{A.15})$$

Next, differentiate S_t with respect to L_t :

$$\frac{\partial S_t}{\partial L_t} = \frac{\partial \Pi_t}{\partial L_t} + \mathbb{F}_t \left[M_{t+1} \frac{\partial S_{t+1}}{\partial L_{t+1}} \frac{\partial L_{t+1}}{\partial L_t} \right] \quad (\text{A.16})$$

Substitute $\frac{\partial L_{t+1}}{\partial L_t} = (1 - \delta_t)$ from the employment accumulation equation (A.13):

$$\frac{\partial S_t}{\partial L_t} = \frac{\partial \Pi_t}{\partial L_t} + (1 - \delta_t) \mathbb{F}_t \left[M_{t+1} \frac{\partial S_{t+1}}{\partial L_{t+1}} \right] \quad (\text{A.17})$$

Substitute equation (A.17) for period $t + 1$ into equation (A.15):

$$\frac{\kappa}{q_t} = \mathbb{F}_t \left[M_{t+1} \left(\frac{\partial \Pi_{t+1}}{\partial L_{t+1}} + (1 - \delta_{t+1}) \mathbb{F}_{t+1} \left[M_{t+2} \frac{\partial S_{t+2}}{\partial L_{t+2}} \right] \right) \right] \quad (\text{A.18})$$

Finally, substitute in (A.15) for period $t + 1$ to arrive at the *hiring equation*:

$$\underbrace{\frac{\kappa}{q_t}}_{\text{Cost of hiring}} = \underbrace{\mathbb{F}_t \left[M_{t+1} \left(\pi_{t+1} + (1 - \delta_{t+1}) \frac{\kappa}{q_{t+1}} \right) \right]}_{\text{Expected discounted value of hiring}} \quad (\text{A.19})$$

where $\pi_t \equiv \frac{\partial \Pi_t}{\partial L_t}$ is the profit flow from the marginal hired worker.

The hiring equation relates the marginal cost of hiring $\frac{\kappa}{q_t}$ with the expected marginal value of hiring to the firm, which equals the future expected marginal benefits of hiring discounted to present value with the stochastic discount factor M_{t+1} . The future marginal benefits of hiring include π_{t+1} , the future marginal product of labor net of the wage rate, plus the future marginal value of hiring, which equals the future marginal cost of hiring $\frac{\kappa}{q_{t+1}}$ net of separation $(1 - \delta_{t+1})$. During recessions, job filling rates q_t are high, which makes the cost of hiring κ/q_t low. The low cost of hiring must be rationalized by either low expected discounted profit flows $\mathbb{F}_t[M_{t+1}\pi_{t+1}]$ or low future value of hiring $(1 - \delta_{t+1})\frac{\kappa}{q_{t+1}}$.

The hiring equation is the labor market analogue of the optimality condition for physical capital in the q theory of investment (Hayashi, 1982), where κ/q_t is the upfront cost of investment analogous to Tobin's marginal q and δ_{t+1} is the depreciation rate.

Stock Price Next, I derive the firm's stock price implied by the optimal hiring decision (Liu et al. (2009), Belo et al. (2023)). Assume a constant returns to scale (CRS) production function so that marginal profits equal average profits:

$$\pi_{t+1} L_{t+1} = \frac{\partial \Pi_{t+1}}{\partial L_{t+1}} L_{t+1} = \Pi_{t+1} \quad (\text{A.20})$$

Multiply both sides of the hiring equation by the number of employees L_{t+1} :

$$\frac{\kappa}{q_t} L_{t+1} = \mathbb{F}_t \left[M_{t+1} \left(\pi_{t+1} L_{t+1} + (1 - \delta_{t+1}) \frac{\kappa}{q_{t+1}} L_{t+1} \right) \right] \quad (\text{A.21})$$

Substitute in the employment accumulation equation (A.13) and rearrange terms:

$$\frac{\kappa}{q_t} L_{t+1} = \mathbb{F}_t \left[M_{t+1} \left(\pi_{t+1} L_{t+1} + \frac{\kappa}{q_{t+1}} (L_{t+2} - q_{t+1} V_{t+1}) \right) \right] \quad (\text{A.22})$$

$$= \mathbb{F}_t \left[M_{t+1} \left(\pi_{t+1} L_{t+1} - \kappa V_{t+1} + \frac{\kappa}{q_{t+1}} L_{t+2} \right) \right] \quad (\text{A.23})$$

Use the constant returns to scale assumption to simplify $\pi_{t+1} L_{t+1} - \kappa V_{t+1} = \Pi_{t+1} - \kappa V_{t+1} = E_{t+1}$:

$$\frac{\kappa}{q_t} L_{t+1} = \mathbb{F}_t \left[M_{t+1} \left(E_{t+1} + \frac{\kappa}{q_{t+1}} L_{t+2} \right) \right] \quad (\text{A.24})$$

Substitute the equation recursively and apply the law of iterated expectations for $\mathbb{F}_t[\cdot]$:

$$\frac{\kappa}{q_t} L_{t+1} = \mathbb{F}_t \left[\sum_{j=1}^{\infty} M_{t,t+j} E_{t+j} \right] + \lim_{T \rightarrow \infty} \mathbb{F}_t \left[M_{t,t+T} \frac{\kappa}{q_{t+T}} L_{t+T+1} \right] \quad (\text{A.25})$$

The first term on the right-hand side is the firm's stock price $P_t \equiv S_t - E_t$, which is the firm's ex-dividend equity value. Take the second term to zero by applying a transversality condition to arrive at an equation that relates the total cost of hiring with the firm's stock price:

$$\frac{\kappa}{q_t} L_{t+1} = P_t \quad (\text{A.26})$$

where employment L_{t+1} is determined at the end of date t under our timing convention from equation (A.9).

Log-linear Approximation Take logarithms of both sides of the firm's stock price equation (A.26) and rearrange terms:

$$\log \kappa - \log q_t = \log \frac{P_t}{L_{t+1}} \quad (\text{A.27})$$

$$= \log \frac{P_t}{E_t} - \log \frac{E_{t+1}}{E_t} + \log \frac{E_{t+1}}{L_{t+1}} \quad (\text{A.28})$$

$$\equiv pe_t - \Delta e_{t+1} + el_{t+1} \quad (\text{A.29})$$

where I define $pe_t \equiv \log \frac{P_t}{E_t}$, $\Delta e_{t+1} \equiv \log \frac{E_{t+1}}{E_t}$, and $el_{t+1} \equiv \log \frac{E_{t+1}}{L_{t+1}}$ for notational convenience. To express the price-earnings ratio pe_t in terms of forward-looking variables, start by log-linearizing the price-dividend ratio $pd_t = \log(P_t/D_t)$ around its long-term average \overline{pd} (Campbell and Shiller, 1988):

$$pd_t = c_{pd} + \Delta d_{t+1} - r_{t+1} + \rho pd_{t+1} \quad (\text{A.30})$$

where c_{pd} is a linearization constant, $r_{t+1} \equiv \log(\frac{P_{t+1}+D_{t+1}}{P_t})$ is the log stock return (with dividends), and $\rho \equiv \frac{\exp(\overline{pd})}{1+\exp(\overline{pd})} = 0.98$ is a persistence parameter that arises from the log linearization. Rewrite

the equation in terms of log price-earnings instead of log price-dividends by using the identity $pe_t = pd_t + de_t$, where de_t log payout ratio:

$$pe_t = c_{pd} + \Delta e_{t+1} - r_{t+1} + \rho pe_{t+1} + (1 - \rho) de_{t+1} \quad (\text{A.31})$$

Since $1 - \rho \approx 0$ and the payout ratio de_t is bounded, $(1 - \rho) de_{t+1}$ can be approximated as a constant, i.e., $c_{pe} \approx c_{pd} + (1 - \rho) de_{t+1}$ (De La O et al., 2024):

$$pe_t \approx c_{pe} + \Delta e_{t+1} - r_{t+1} + \rho pe_{t+1} \quad (\text{A.32})$$

Recursively substitute for the next h periods

$$pe_t = \sum_{j=1}^h \rho^{j-1} (c_{pe} + \Delta e_{t+j} - r_{t+j}) + \rho^h pe_{t+h} \quad (\text{A.33})$$

Decomposition of Job Filling Rates Substitute the log-linearized price-earnings ratio in equation (A.33) into the log stock price in equation (A.29):

$$\log q_t = \log \kappa - pe_t + \Delta e_{t+1} - el_{t+1} \quad (\text{A.34})$$

$$= \log \kappa - \left[\sum_{j=1}^h \rho^{j-1} (c_{pe} + \Delta e_{t+j} - r_{t+j}) + \rho^h pe_{t+h} \right] + \Delta e_{t+1} - el_{t+1} \quad (\text{A.35})$$

Rearrange and collect terms to obtain an ex-post decomposition of the job filling rate:

$$\log q_t = c_q + \underbrace{\sum_{j=1}^h \rho^{j-1} r_{t+j}}_{r_{t,t+h}} - \underbrace{\left[el_{t+1} + \sum_{j=2}^h \rho^{j-1} \Delta e_{t+j} \right]}_{e_{t,t+h}} - \underbrace{\rho^h pe_{t+h}}_{pe_{t,t+h}} \quad (\text{A.36})$$

where $c_q \equiv \log \kappa - \frac{c_{pe}(1-\rho^h)}{1-\rho}$ is a constant. The equation decomposes the job filling rate into future discount rates $r_{t,t+h} \equiv \sum_{j=1}^h \rho^{j-1} r_{t+j}$, cash flows $e_{t,t+h} \equiv el_{t+1} + \sum_{j=2}^h \rho^{j-1} \Delta e_{t+j}$, and price-earnings $pe_{t,t+h} \equiv \rho^h pe_{t+h}$. The cash flow component consists of one period ahead log earnings-employment el_{t+1} , which captures news about short-term cash flows, and $j = 2, \dots, h$ period ahead log earnings growth Δe_{t+j} , which captures long-term news about cash flows.¹ $pe_{t,t+h}$ is a terminal value that captures other long-term influences beyond h periods into the future not already captured in discount rates and cash flows.

Since equation (A.36) holds both ex-ante and ex-post, it can be evaluated under either subjective or rational expectations. The *subjective decomposition* replaces ex-post realizations of future outcomes with their subjective expectations $\mathbb{F}_t[\cdot]$:

$$\log q_t = c_q + \underbrace{\sum_{j=1}^h \rho^{j-1} \mathbb{F}_t[r_{t+j}]}_{\mathbb{F}_t[r_{t,t+h}]} - \underbrace{\left[\mathbb{F}_t[el_{t+1}] + \sum_{j=2}^h \rho^{j-1} \mathbb{F}_t[\Delta e_{t+j}] \right]}_{\mathbb{F}_t[e_{t,t+h}]} - \underbrace{\rho^h \mathbb{F}_t[pe_{t+h}]}_{\mathbb{F}_t[pe_{t,t+h}]} \quad (\text{A.37})$$

¹The earnings-employment ratio can be interpreted as a measure of the marginal product of labor under constant returns to scale (David et al., 2022).

Alternatively, the *rational decomposition* replaces ex-post realizations of future outcomes with their rational expectations $\mathbb{E}_t[\cdot]$:

$$\log q_t = c_q + \underbrace{\sum_{j=1}^h \rho^{j-1} \mathbb{E}_t[r_{t+j}]}_{\mathbb{E}_t[r_{t,t+h}]} - \underbrace{\left[\mathbb{E}_t[el_{t+1}] + \sum_{j=2}^h \rho^{j-1} \mathbb{E}_t[\Delta e_{t+j}] \right]}_{\mathbb{E}_t[e_{t,t+h}]} - \underbrace{\rho^h \mathbb{E}_t[pe_{t+h}]}_{\mathbb{E}_t[pe_{t,t+h}]} \quad (\text{A.38})$$

Comparing these decompositions can quantify how belief distortions affect the job filling rate.

Estimation The econometrician can estimate the variance decomposition using predictive regressions of each expected outcome on the current job filling rate. For the subjective decomposition, demean each variable in equation (A.37), multiply both sides by the current log job filling rate $\log q_t$, and take the sample average:

$$Var[\log q_t] = Cov[\mathbb{F}_t[r_{t,t+h}], \log q_t] - Cov[\mathbb{F}_t[e_{t,t+h}], \log q_t] - Cov[\mathbb{F}_t[pe_{t,t+h}], \log q_t] \quad (\text{A.39})$$

where $Var[\cdot]$ and $Cov[\cdot]$ are sample variances and covariances based on data observed over a historical sample. Finally, divide both sides by $Var[\log q_t]$ to decompose its variance:

$$1 = \underbrace{\frac{Cov[\mathbb{F}_t[r_{t,t+h}], \log q_t]}{Var[\log q_t]}}_{\text{Discount Rate News}} - \underbrace{\frac{Cov[\mathbb{F}_t[e_{t,t+h}], \log q_t]}{Var[\log q_t]}}_{\text{Cash Flow News}} - \underbrace{\frac{Cov[\mathbb{F}_t[pe_{t,t+h}], \log q_t]}{Var[\log q_t]}}_{\text{Future Price-Earnings News}} \quad (\text{A.40})$$

The left-hand side represents the full variability in job filling rates, hence is equal to one. Each term on the right reflects the share explained by subjective expectations of discount rates, cash flows, or price-earnings ratios. Under stationarity, the econometrician can estimate these shares using the OLS coefficients from regressing $\mathbb{F}_t[r_{t,t+h}]$, $\mathbb{F}_t[e_{t,t+h}]$, and $\mathbb{F}_t[pe_{t,t+h}]$ on the current log job filling rate $\log q_t$, respectively.

Finally, the decomposition under rational expectations can be estimated similarly based on equation (A.38) by replacing the subjective expectation $\mathbb{F}_t[\cdot]$ with its rational counterpart $\mathbb{E}_t[\cdot]$:

$$1 = \underbrace{\frac{Cov[\mathbb{E}_t[r_{t,t+h}], \log q_t]}{Var[\log q_t]}}_{\text{Discount Rate News}} - \underbrace{\frac{Cov[\mathbb{E}_t[e_{t,t+h}], \log q_t]}{Var[\log q_t]}}_{\text{Cash Flow News}} - \underbrace{\frac{Cov[\mathbb{E}_t[pe_{t,t+h}], \log q_t]}{Var[\log q_t]}}_{\text{Future Price-Earnings News}} \quad (\text{A.41})$$

Under stationarity, the econometrician can estimate these shares using the OLS coefficients from regressing $\mathbb{E}_t[r_{t,t+h}]$, $\mathbb{E}_t[e_{t,t+h}]$, and $\mathbb{E}_t[pe_{t,t+h}]$ on the current log job filling rate $\log q_t$, respectively.

C Time-Series Decomposition of the Unemployment Rate

I derive a decomposition of the unemployment rate by log linearizing the unemployment accumulation equation from equation (2):

$$U_{t+1} = \delta_t(1 - U_t) + (1 - q_t\theta_t)U_t \quad (\text{A.42})$$

Denote the steady state values without time subscripts: U , δ , q , and θ . In steady state:

$$U = \delta(1 - U) + (1 - q\theta)U \quad (\text{A.43})$$

Define log deviations from steady state as $\hat{X}_t = \log(X_t) - \log(X)$ for some variable X . Log-linearizing the accumulation equation around the steady state involves taking a first-order Taylor approximation:

$$Ue^{\hat{U}_{t+1}} \approx \delta e^{\hat{\delta}_t}(1 - Ue^{\hat{U}_t}) + (1 - q\theta e^{\hat{q}_t + \hat{\theta}_t})Ue^{\hat{U}_t} \quad (\text{A.44})$$

Use the approximation $Xe^{x_t} \approx X(1 + x_t)$, expand, and simplify:

$$U(1 + \hat{U}_{t+1}) \approx \delta(1 + \hat{\delta}_t)(1 - U(1 + \hat{U}_t)) + (1 - q\theta(1 + \hat{q}_t + \hat{\theta}_t))U(1 + \hat{U}_t) \quad (\text{A.45})$$

or equivalently

$$U + U\hat{U}_{t+1} \approx \delta(1 + \hat{\delta}_t)(1 - U - U\hat{U}_t) + (1 - q\theta - q\theta\hat{q}_t - q\theta\hat{\theta}_t)U(1 + \hat{U}_t) \quad (\text{A.46})$$

Use the steady state equation and collect terms with log deviations:

$$U\hat{U}_{t+1} \approx \delta(1 - U)\hat{\delta}_t - \delta U\hat{U}_t - q\theta U\hat{q}_t - q\theta U\hat{\theta}_t + U(1 - q\theta)\hat{U}_t \quad (\text{A.47})$$

Divide both sides by U :

$$\hat{U}_{t+1} \approx \frac{\delta(1 - U)}{U}\hat{\delta}_t - \delta\hat{U}_t - q\theta\hat{q}_t - q\theta\hat{\theta}_t + (1 - q\theta)\hat{U}_t \quad (\text{A.48})$$

or equivalently:

$$\hat{U}_{t+1} \approx \frac{\delta(1 - U)}{U}\hat{\delta}_t + (1 - \delta - q\theta)\hat{U}_t - q\theta\hat{q}_t - q\theta\hat{\theta}_t \quad (\text{A.49})$$

The steady state relationship $\delta(1 - U) = q\theta U$ implies:

$$\frac{\delta(1 - U)}{U} = q\theta \quad (\text{A.50})$$

Substitute this back into our equation:

$$\hat{U}_{t+1} \approx q\theta\hat{\delta}_t + (1 - \delta - q\theta)\hat{U}_t - q\theta\hat{q}_t - q\theta\hat{\theta}_t \quad (\text{A.51})$$

The result is a linear relationship between a decomposition of the log job filling rate and the log-linearized unemployment accumulation equation.

Finally, substitute in equation (11), which is a decomposition of the job filling rate \hat{q}_t into discount rate, cash flow, and future price-earnings components:

$$\hat{U}_{t+1} \approx - \underbrace{q\theta \cdot \hat{r}_{t,t+h}}_{\text{Discount Rate}} + \underbrace{q\theta \cdot \hat{e}_{t,t+h}}_{\text{Cash Flow}} + \underbrace{q\theta \cdot \hat{p}e_{t,t+h}}_{\text{Future Price-Earning}} + \underbrace{(1 - \delta - q\theta) \cdot \hat{U}_t - q\theta \cdot \hat{\theta}_t + q\theta \cdot \hat{\delta}_t}_{\text{Lag Unemployment, Tightness, Separations}} \quad (\text{A.52})$$

The equation holds both ex-ante and ex-post. Therefore, I compare results from evaluating the equation under subjective $\mathbb{F}_t[\cdot]$ or rational $\mathbb{E}_t[\cdot]$ expectations. The decomposition can be estimated using regressions of the log unemployment rate on each of the components shown in the equation. To make sure all variables in the regression are stationary, I estimate the regressions in first differences.

C.1 Regional Model

Model Setup This section presents a multi-area, multi-sector search-and-matching model with imperfect mobility across sectors, building from the models in Kehoe et al. (2019) and Chodorow-Reich and Wieland (2020). The economy consists of a continuum of islands indexed by s . Each island produces a differentiated variety of tradable goods that is consumed everywhere and a nontradable good. Both of these goods are produced using intermediate goods. Each consumer is endowed with one of two types of skills which are used in different intensities in the nontradable and tradable goods sectors. Labor is immobile across islands but can switch sectors.² Consumers receive utility from a composite consumption good that is either purchased in the market or produced at home. Consumers and firms are ex-ante homogeneous and share the same subjective expectation $\mathbb{F}_t[\cdot]$. The islands only differ in the shocks that hit them.

Preferences and demand The composite consumption good on island s is produced from nontradable goods $X_{s,N,t}$ and tradable goods $X_{s,T,t}$

$$X_{s,t} = \left[\tau^{\frac{1}{\mu}} (X_{s,N,t})^{1-\frac{1}{\mu}} + (1-\tau)^{\frac{1}{\mu}} (X_{s,T,t})^{1-\frac{1}{\mu}} \right]^{\frac{\mu}{\mu-1}} \quad (\text{A.53})$$

where μ is the elasticity of substitution between tradable and nontradable goods. The demand for nontradable and tradable goods on island s is

$$X_{s,N,t} = \tau \left(\frac{P_{s,N,t}}{P_{s,t}} \right)^{-\mu} X_{s,t}, \quad X_{s,T,t} = (1-\tau) \left(\frac{P_{s,T,t}}{P_{s,t}} \right)^{-\mu} X_{s,t} \quad (\text{A.54})$$

where $P_{s,N,t}$ is the price of the nontradable good and $P_{s,T,t}$ is the world price of the composite tradable good. The price of the composite consumption good on island s is

$$P_{s,t} = [\tau P_{s,N,t}^{1-\mu} + (1-\tau) P_{s,T,t}^{1-\mu}]^{\frac{1}{1-\mu}} \quad (\text{A.55})$$

The tradable good is a composite of varieties of differentiated tradable goods produced in all islands s'

$$X_{s,T,t} = \left[\int X_{s,T,t,s'}^{\frac{\mu_T-1}{\mu_T}} ds' \right]^{\frac{\mu_T}{\mu_T-1}} \quad (\text{A.56})$$

where $X_{s,T,t,s'}$ is the amount of the variety of tradable good produced on island s' and consumed on island s . μ_T is the elasticity of substitution between varieties produced on different islands. Let $P_{s',T,t}$ be the price of tradable variety produced on island s' . Assume that there are no costs of shipping goods from one island to another, so that the law of one price holds and all islands purchase the variety s at the common price $P_{s,T,t}$. The price of the composite tradable good is common to all islands

$$P_{T,t} = \left[\int P_{s,T,t}^{1-\mu_T} ds \right]^{\frac{1}{1-\mu_T}} \quad (\text{A.57})$$

²This assumption aligns with empirical evidence indicating that labor markets are predominantly regional in nature (Manning and Petrongolo, 2017).

The demand on island s' for a tradable variety produced on island s is therefore

$$X_{s',T,t,s} = \left[\frac{P_{s,T,t}}{P_{T,t}} \right]^{-\mu_T} X_{s',T,t} \quad (\text{A.58})$$

so that the world demand for tradable goods produced by island s is

$$Y_{s,T,t} = \int X_{s',T,t,s} ds' = \left[\frac{P_{s,T,t}}{P_{T,t}} \right]^{-\mu_T} Y_{T,t} \quad (\text{A.59})$$

where $Y_{T,t} = \int X_{s',T,t} ds'$ is the world demand for the composite tradable good. Since any individual island is of measure zero, shocks to an individual island do not affect either the world aggregate price of tradables $P_{T,t}$ or the world demand for tradables $Y_{T,t}$. Normalize the constant world price of the composite tradable good $P_{T,t}$ to one so that the composite tradable good is the numeraire.

Family's problem Each family of workers on island s chooses sequences for consumption $\{C_{s,t}\}$ and assets $\{A_{s,t+1}\}$ to maximize the present discounted value of consumption

$$\max_{C_{s,t}, A_{s,t+1}} \sum_{t=0}^{\infty} \beta^t u(C_{s,t}) \quad (\text{A.60})$$

where the family's consumption $C_{s,t} = X_{s,t} + b_{s,t}$ is the sum of goods purchased in the market $X_{s,t}$, and produced at home $b_{s,t}$ which can be consumed only by that family. The budget constraint is

$$P_{s,t} X_{s,t} + q^A A_{s,t+1} = Y_{s,t} + E_{s,t} + A_{s,t} \quad (\text{A.61})$$

where $P_{s,t}$ is the price of the composite consumption good on the island, $A_{s,t}$ are the family's assets, and the family saves or borrows at a constant world bond price $q^A > \beta$. $Y_{s,t}$ is the income of the family's workers in the form of wages

$$Y_{s,t} \equiv \sum_{i \in \{\mathcal{N}, \mathcal{T}\}} w_{s,i,t} L_{s,i,t} \quad (\text{A.62})$$

and $E_{s,t}$ are profits from the firms the family owns on island s

$$E_{s,t} \equiv \sum_{i \in \{\mathcal{N}, \mathcal{T}\}} E_{s,i,t} = \sum_{i \in \{\mathcal{N}, \mathcal{T}\}} [(z_{s,i,t} - w_{s,i,t}) L_{s,i,t} - \kappa V_{s,i,t}] \quad (\text{A.63})$$

where $z_{s,i,t}$ is a sectoral labor productivity shock, $w_{s,i,t}$ is the wage of an employed worker, $L_{s,i,t}$ is the measure of employed workers, and $V_{s,i,t}$ is the measure of vacancies for producing intermediate goods of type i on island s .

From the first-order condition for the family's problem, we can derive the shadow price of the composite consumption good at date t in units of the composite consumption good at date 0 on island s as

$$M_{s,t} = \beta^t \frac{u'(C_{s,t})/P_{s,t}}{u'(C_{s,0})/P_{s,0}} \quad (\text{A.64})$$

Technology Nontradable and tradable goods are produced with locally produced intermediate goods. These intermediate goods are used by the nontradable and tradable sectors in different proportions. This setup effectively introduces costs of sectoral reallocations of workers because it implies a curved production possibility frontier between nontradable and tradable goods.

The economy has two types of intermediate goods: Type \mathcal{N} and type \mathcal{T} goods. The technology for producing nontradable goods disproportionately uses type \mathcal{N} goods, whereas the technology for producing tradable goods disproportionately uses type \mathcal{T} goods according to the production technologies

$$Y_{s,N,t} = A(Y_{s,N,t}^{\mathcal{N}})^{\nu}(Y_{s,N,t}^{\mathcal{T}})^{1-\nu}, \quad Y_{s,T,t} = A(Y_{s,T,t}^{\mathcal{N}})^{1-\nu}(Y_{s,T,t}^{\mathcal{T}})^{\nu} \quad (\text{A.65})$$

with $\nu \geq 1/2$. $Y_{s,N,t}^{\mathcal{N}}$ and $Y_{s,T,t}^{\mathcal{N}}$ denote the use of intermediate inputs of type \mathcal{N} in the production of nontradable and tradable goods, whereas $Y_{s,N,t}^{\mathcal{T}}$ and $Y_{s,T,t}^{\mathcal{T}}$ denote the use of intermediate inputs of type \mathcal{T} in the production of nontradable and tradable goods. Both nontradable goods producers and tradable goods producers are competitive and take as given the price of their goods, $P_{s,N,t}$ and $P_{s,T,t}$. The demands for intermediate inputs in the nontradable goods sector are

$$Y_{s,N,t}^{\mathcal{N}} = \nu \left(\frac{P_{s,t}^{\mathcal{T}}}{P_{s,t}^{\mathcal{N}}} \right)^{1-\nu} Y_{s,N,t}, \quad Y_{s,N,t}^{\mathcal{T}} = (1-\nu) \left(\frac{P_{s,t}^{\mathcal{N}}}{P_{s,t}^{\mathcal{T}}} \right)^{\nu} Y_{s,N,t} \quad (\text{A.66})$$

where $P_{s,t}^{\mathcal{N}}$ and $P_{s,t}^{\mathcal{T}}$ are prices of the intermediate goods \mathcal{N} and \mathcal{T} . The equation was derived under the normalization $A = \nu^{-\nu}(1-\nu)^{1-\nu}$. Likewise, the demands for intermediate inputs in the tradable goods sector are

$$Y_{s,T,t}^{\mathcal{N}} = (1-\nu) \left(\frac{P_{s,t}^{\mathcal{N}}}{P_{s,t}^{\mathcal{T}}} \right)^{\nu} Y_{s,T,t}, \quad Y_{s,T,t}^{\mathcal{T}} = \nu \left(\frac{P_{s,t}^{\mathcal{T}}}{P_{s,t}^{\mathcal{N}}} \right)^{1-\nu} Y_{s,T,t} \quad (\text{A.67})$$

Adding up the demands for each type of intermediate good by the two sectors gives the total demand on island s for intermediate goods of type i

$$Y_{s,t}^i = Y_{s,N,t}^i + Y_{s,T,t}^i \quad (\text{A.68})$$

Production of these intermediate goods is given by

$$Y_{s,t}^i = z_{s,i,t} \cdot L_{s,i,t} \quad (\text{A.69})$$

where $L_{s,i,t}$ is the measure of employed workers producing intermediate goods of type i on island s . $z_{s,i,t}$ represents exogenous labor productivity for producing intermediate goods of type i on island s . Zero profit conditions in nontradable and tradable goods sectors imply

$$P_{s,N,t} = (P_{s,t}^{\mathcal{N}})^{\nu}(P_{s,t}^{\mathcal{T}})^{1-\nu}, \quad P_{s,T,t} = (P_{s,t}^{\mathcal{N}})^{1-\nu}(P_{s,t}^{\mathcal{T}})^{\nu} \quad (\text{A.70})$$

Assume that there are measures of consumers $\pi^{\mathcal{N}}$ and $\pi^{\mathcal{T}} = 1 - \pi^{\mathcal{N}}$ in occupations \mathcal{N} and \mathcal{T} who supply labor to produce the two types of intermediate goods \mathcal{N} and \mathcal{T} , respectively. Consumers

in occupation \mathcal{N} can produce good \mathcal{N} , and consumers in occupation \mathcal{T} can produce good \mathcal{T} . Consumers are hired by intermediate goods firms that produce intermediate goods of either type \mathcal{N} or \mathcal{T} . These goods are then sold at competitive prices $P_{s,t}^{\mathcal{N}}$ and $P_{s,t}^{\mathcal{T}}$ to firms in the nontradable and tradable goods sectors.³

This setup captures in a simple way the idea that switching sectors is relatively easy, whereas switching occupations is difficult. Note that any individual consumer faces no cost of switching sectors. But if a positive measure of consumer moves from one sector to the other, then the marginal revenue product of the consumers in the new sector falls and so do wages. This reduction in marginal revenue products acts like a switching cost in the aggregate.

Labor market Firms that produce intermediate good $i \in \{\mathcal{N}, \mathcal{T}\}$ post vacancies for consumers in occupation i , who produce intermediate good i when matched. Assume that consumers cannot switch occupations, so the measure of consumers in each occupation is fixed. The values of consumers in occupation i of island s are

$$\widetilde{W}_{s,i,t}(z_{s,i,t}) = w_{s,i,t}(z_{s,i,t}) \quad (\text{A.71})$$

$$+ (1 - \delta)\mathbb{F}_t[M_{s,t,t+1}\widetilde{W}_{s,i,t+1}(z_{s,i,t+1})] \quad (\text{A.72})$$

$$+ \delta\mathbb{F}_t[M_{s,t,t+1}\widetilde{U}_{s,i,t+1}(z_{s,i,t+1})] \quad (\text{A.73})$$

for employed consumers, and

$$\widetilde{U}_{s,i,t}(z_{s,i,t}) = P_t b_{s,t}(z_{s,i,t}) \quad (\text{A.74})$$

$$+ \mathbb{F}_t[M_{s,t,t+1}f_{s,i,t}(z_{s,i,t})\widetilde{W}_{s,i,t+1}(z_{s,i,t+1})] \quad (\text{A.75})$$

$$+ \mathbb{F}_t[M_{s,t,t+1}(1 - f_{s,i,t}(z_{s,i,t}))\widetilde{U}_{s,i,t+1}(z_{s,i,t+1})] \quad (\text{A.76})$$

for nonemployed consumers. $w_{s,i,t}(z_{s,i,t})$ is the wage received by a consumer in occupation i . $f_{s,i,t}(z_{s,i,t})$ is the job-finding probability of a consumer in occupation i . $b_{s,t}(z_{s,i,t})$ is the output of a consumer when not employed. δ is an exogenous separation probability. Subjective expectations $\mathbb{F}_t[\cdot]$ are with respect to next period's productivity $z_{s,i,t+1}$.

The value of a firm producing intermediate good i matched with a consumer in occupation i with productivity $z_{s,i,t}$ is

$$\widetilde{J}_{s,i,t}(z_{s,i,t}) = P_{s,i,t}z_{s,i,t} - w_{s,i,t}(z_{s,i,t}) + (1 - \delta)\mathbb{F}_t[M_{s,t,t+1}\widetilde{J}_{s,i,t+1}(z_{s,i,t+1})] \quad (\text{A.77})$$

At date t a consumer in occupation i matched with a firm in intermediate good sector i produces $z_{s,i,t}$ units of good i , which sells for $P_{s,i,t}z_{s,i,t}$, and the firm pays the consumer $w_{s,i,t}(z_{s,i,t})$. The

³It is equivalent to think that the consumers in each occupation work in the sector that purchases the goods they produce. Under this interpretation, we can think of consumers in occupation \mathcal{T} as being employed in sectors N and T and consumers in occupation \mathcal{N} as also being employed in sectors N and T in different proportions. Sector N employs consumers in occupation \mathcal{N} relatively more intensively, whereas sector T employs consumers in occupation \mathcal{T} relatively more intensively.

cost of posting a vacancy is κ units of the composite tradable good. Free entry for intermediate goods producers in the labor market for workers in occupation i implies

$$\kappa = q_{s,i,t}(z_{s,i,t}) \cdot \mathbb{E}_t[M_{s,t,t+1} \tilde{J}_{s,i,t+1}(z_{s,i,t+1})] \quad (\text{A.78})$$

The matches of firms that produce intermediate good i with consumers are

$$L_{s,i,t}(z_{s,i,t}) = \frac{U_{s,i,t}(z_{s,i,t})V_{s,i,t}(z_{s,i,t})}{[U_{s,i,t}(z_{s,i,t})^\eta + V_{s,i,t}(z_{s,i,t})^\eta]^{1/\eta}} \quad (\text{A.79})$$

where $U_{s,i,t}(z_{s,i,t})$ is the measure of nonemployed consumers and $V_{s,i,t}(z_{s,i,t})$ is the measure of posted vacancies to attract such consumers. The parameter η governs the sensitivity of $f_{s,i,t}(z_{s,i,t})$ to $\theta_{s,i,t}$. The worker job-finding rate $f_{s,i,t}(z_{s,i,t})$ and firm job-filling rate $q_{s,i,t}(z_{s,i,t})$ are

$$f_{s,i,t}(z_{s,i,t}) = \frac{L_{s,i,t}(z_{s,i,t})}{U_{s,i,t}(z_{s,i,t})} = \frac{\theta_{s,i,t}(z_{s,i,t})}{[1 + \theta_{s,i,t}(z_{s,i,t})^\eta]^{1/\eta}} \quad (\text{A.80})$$

$$q_{s,i,t}(z_{s,i,t}) = \frac{L_{s,i,t}(z_{s,i,t})}{V_{s,i,t}(z_{s,i,t})} = \frac{1}{[1 + \theta_{s,i,t}(z_{s,i,t})^\eta]^{1/\eta}} \quad (\text{A.81})$$

where $\theta_{s,i,t}(z_{s,i,t}) = V_{s,i,t}(z_{s,i,t})/U_{s,i,t}(z_{s,i,t})$ is the vacancy to nonemployment ratio.

The Nash bargaining problem determines the wage $w_{s,i,t}(z_{s,i,t})$ in any given match

$$\max_w [\tilde{W}_{s,i,t}(z_{s,i,t}) - \tilde{U}_{s,i,t}(z_{s,i,t})]^\gamma \tilde{J}_{s,i,t}(z_{s,i,t})^{1-\gamma} \quad (\text{A.82})$$

where γ is a consumer's bargaining weight. Defining the surplus of a match between a firm and a consumer as $\tilde{S}_{s,i,t}(z_{s,i,t}) = \tilde{W}_{s,i,t}(z_{s,i,t}) - \tilde{U}_{s,i,t}(z_{s,i,t}) + \tilde{J}_{s,i,t}(z_{s,i,t})$, Nash bargaining implies that firms and consumers split this surplus according to

$$\tilde{W}_{s,i,t}(z_{s,i,t}) - \tilde{U}_{s,i,t}(z_{s,i,t}) = \gamma \tilde{S}_{s,i,t}(z_{s,i,t}), \quad \tilde{J}_{s,i,t}(z_{s,i,t}) = (1 - \gamma) \tilde{S}_{s,i,t}(z_{s,i,t}) \quad (\text{A.83})$$

Equilibrium Consider now the market-clearing conditions. Market clearing for the two types of intermediate goods requires that

$$Y_{s,i,t} = z_{s,i,t} \cdot L_{s,i,t} = Y_{s,N,t}^i + Y_{s,T,t}^i, \quad i \in \{\mathcal{N}, \mathcal{T}\} \quad (\text{A.84})$$

the left side of this equation is the total amount of intermediate goods of type i produced by employed workers in occupation i on island s , $L_{s,i,t}$. The right side is the total amount of these intermediate goods used by firms in the nontradable and tradable goods sectors on that island. Employment in the nontradable goods sector on island s is

$$\frac{Y_{s,N,t}^{\mathcal{N}}}{Y_{s,t}^{\mathcal{N}}} L_{s,\mathcal{N},t} + \frac{Y_{s,N,t}^{\mathcal{T}}}{Y_{s,t}^{\mathcal{T}}} L_{s,\mathcal{T},t} \quad (\text{A.85})$$

Employment in the tradable goods sector on island s is

$$\frac{Y_{s,T,t}^{\mathcal{N}}}{Y_{s,t}^{\mathcal{N}}} L_{s,\mathcal{N},t} + \frac{Y_{s,T,t}^{\mathcal{T}}}{Y_{s,t}^{\mathcal{T}}} L_{s,\mathcal{T},t} \quad (\text{A.86})$$

The relative demand effect on employment in the two sectors captures the idea that, since $Y_{s,N,t}^i/Y_{s,i,t} + Y_{s,T,t}^i/Y_{s,i,t} = 1$ for $i \in \{\mathcal{N}, \mathcal{T}\}$, any shift in demand from the nontradable goods sector on an island, holding fixed total employment on the island, decreases employment in the nontradable goods sector and increases it in the tradable goods sector on the island.

Market clearing for nontradable goods requires that the demand for nontradable goods on island s equal the amount on nontradable goods produced on island s

$$X_{s,N,t} = A(Y_{s,N,t}^{\mathcal{N}})^{\nu} (Y_{s,N,t}^{\mathcal{T}})^{1-\nu} \quad (\text{A.87})$$

Similarly, market clearing for tradable goods requires that the world demand for tradable goods produced from island s equal the amount of tradable goods produced on island s

$$Y_{s,T,t} = A(Y_{s,T,t}^{\mathcal{N}})^{1-\nu} (Y_{s,T,t}^{\mathcal{T}})^{\nu} \quad (\text{A.88})$$

Decomposition of Regional Job Filling Rates Combine the value of the worker to the firm with the zero-profit condition for entering firms, substitute recursively, and apply the law of iterated expectations

$$\frac{\kappa}{q_{s,i,t}} = \mathbb{F}_t[M_{s,t,t+1} \tilde{J}_{s,i,t+1}] \quad (\text{A.89})$$

$$= \mathbb{F}_t[M_{s,t,t+1} (P_{s,i,t+1} z_{s,i,t+1} - w_{s,i,t+1} + (1-\delta) \mathbb{F}_{t+1}[M_{s,t+1,t+2} \tilde{J}_{s,i,t+2}])] \quad (\text{A.90})$$

$$= \mathbb{F}_t \left[M_{s,t,t+1} \left(P_{s,i,t+1} z_{s,i,t+1} - w_{s,i,t+1} + (1-\delta) \frac{\kappa}{q_{s,i,t+1}} \right) \right] \quad (\text{A.91})$$

Multiply both sides by the level of employment in sector i island s

$$\frac{\kappa}{q_{s,i,t}} L_{s,i,t+1} = \mathbb{F}_t \left[M_{s,t,t+1} \left((P_{s,i,t+1} z_{s,i,t+1} - w_{s,i,t+1}) L_{s,i,t+1} + (1-\delta) \frac{\kappa}{q_{s,i,t+1}} L_{s,i,t+1} \right) \right] \quad (\text{A.92})$$

Substitute in the law of motion for employment $L_{s,i,t+1} = (1-\delta)L_{s,i,t} + q_{s,i,t}V_{s,i,t}$

$$\frac{\kappa}{q_{s,i,t}} L_{s,i,t+1} = \mathbb{F}_t \left[M_{s,t,t+1} \left((P_{s,i,t+1} z_{s,i,t+1} - w_{s,i,t+1}) L_{s,i,t+1} - \kappa V_{s,i,t+1} + \frac{\kappa}{q_{s,i,t+1}} L_{s,i,t+2} \right) \right] \quad (\text{A.93})$$

Define the firm's total earnings as profits after vacancy posting costs $E_{s,i,t} \equiv (P_{s,i,t} z_{s,i,t} - w_{s,i,t}) L_{s,i,t} - \kappa V_{s,i,t}$. Substitute recursively and apply the law of iterated expectations

$$\frac{\kappa}{q_{s,i,t}} L_{s,i,t+1} = \mathbb{F}_t \left[\sum_{j=1}^{T-1} \left(\prod_{k=1}^j M_{s,t+k-1,t+k} \right) E_{s,i,t+j} \right] + \mathbb{F}_t \left[M_{s,t+T-1,t+T} \frac{\kappa}{q_{s,i,t+T}} L_{s,i,t+T+1} \right] \quad (\text{A.94})$$

Take limits as $T \rightarrow \infty$ while applying a transversality condition to rule out bubbles

$$\frac{\kappa}{q_{s,i,t}} L_{s,i,t+1} = \mathbb{F}_t \left[\sum_{j=1}^{\infty} \left(\prod_{k=1}^j M_{s,t+k-1,t+k} \right) E_{s,i,t+j} \right] \equiv P_{s,i,t}^E \quad (\text{A.95})$$

where $P_{s,i,t}^E$ is the firm's ex-dividend equity value. Take logarithms on both sides and expand the price-employment ratio $P_{s,i,t}^E/L_{s,i,t+1}$

$$\log q_{s,i,t} = \log \kappa - \log(P_{s,i,t}^E/L_{s,i,t+1}) \quad (\text{A.96})$$

$$= \log \kappa - \log(P_{s,i,t}^E/E_{s,i,t}) + \log(E_{s,i,t+1}/E_{s,i,t}) - \log(E_{s,i,t+1}/L_{s,i,t+1}) \quad (\text{A.97})$$

$$\equiv \log \kappa - pe_{s,i,t} + \Delta e_{s,i,t+1} - el_{s,i,t+1} \quad (\text{A.98})$$

where $pe_{s,i,t} \equiv \log(P_{s,i,t}^E/E_{s,i,t})$ is the log price-earnings ratio, $\Delta e_{s,i,t+1} \equiv \log(E_{s,i,t+1}/E_{s,i,t})$ is log earnings growth, and $el_{s,i,t+1} \equiv \log(E_{s,i,t+1}/L_{s,i,t+1})$ is the log earnings-employment ratio.

Next, we log-linearize the price-earnings ratio by first log-linearizing the price-dividend ratio $pd_{s,i,t}$ around its long-term average \bar{pd} (Campbell and Shiller, 1988):

$$pd_{s,i,t} = c_{pd} + \Delta d_{s,i,t+1} - r_{s,i,t+1} + \rho pd_{s,i,t+1} \quad (\text{A.99})$$

where c_{pd} is a linearization constant, $r_{s,i,t+1} \equiv \log(\frac{P_{s,i,t+1}^E + D_{s,i,t+1}}{P_{s,i,t}^E})$ is the log stock return (with dividends), and $\rho \equiv \frac{\exp(\bar{pd})}{1 + \exp(\bar{pd})} = 0.98$ is a persistence parameter that arises from the log linearization. Rewrite the equation in terms of log price-earnings instead of log price-dividends by using the identity $pe_{s,i,t} = pd_{s,i,t} + de_{s,i,t}$, where $de_{s,i,t}$ log payout ratio:

$$pe_{s,i,t} = c_{pd} + \Delta e_{s,i,t+1} - r_{s,i,t+1} + \rho pe_{s,i,t+1} + (1 - \rho) de_{s,i,t+1} \quad (\text{A.100})$$

Since $1 - \rho \approx 0$ and the payout ratio $de_{s,i,t}$ is bounded, $(1 - \rho) de_{s,i,t+1}$ can be approximated as a constant, i.e., $c_{pe} \approx c_{pd} + (1 - \rho) de_{s,i,t+1}$ (De La O et al., 2024):

$$pe_{s,i,t} \approx c_{pe} + \Delta e_{s,i,t+1} - r_{s,i,t+1} + \rho pe_{s,i,t+1} \quad (\text{A.101})$$

Recursively substitute for the next h periods

$$pe_{s,i,t} = \sum_{j=1}^h \rho^{j-1} (c_{pe} + \Delta e_{s,i,t+j} - r_{s,i,t+j}) + \rho^h pe_{s,i,t+h} \quad (\text{A.102})$$

Substitute log-linearized price-earnings into the hiring equation to obtain a decomposition of the log job filling rate. In this environment, the intermediate good producer's optimal hiring decision satisfies a regional analog of the aggregate hiring equation (see Section C.1 for details)

$$\frac{\kappa}{q_{s,i,t}} = \mathbb{F}_t \left[M_{s,t,t+1} \left(\pi_{s,i,t+1} + (1 - \delta) \frac{\kappa}{q_{s,i,t+1}} \right) \right] \quad (\text{A.103})$$

for a firm in island s that produces intermediate good i at time t . Then the job filling rate $q_{s,i,t}$ in island i can be decomposed as

$$\log q_{s,i,t} = c_q + \underbrace{\sum_{j=1}^h \rho^{j-1} r_{s,i,t+j}}_{r_{s,i,t,t+h}} - \underbrace{\left[el_{s,i,t+1} + \sum_{j=2}^h \rho^{j-1} \Delta e_{s,i,t+j} \right]}_{e_{s,i,t,t+h}} - \underbrace{\rho^h pe_{s,i,t+h}}_{pe_{s,i,t,t+h}} \quad (\text{A.104})$$

where $c_q \equiv \log \kappa - \frac{c_{pe}(1-\rho^h)}{1-\rho}$ is a constant. $r_{s,i,t,t+h}$ captures news about discount rates. $e_{s,i,t,t+h}$ captures news about cash flows, where the one year ahead log earnings-employment ratio $el_{s,i,t+1}$ captures news about short-term cash flows and $j = 2, \dots, h$ year ahead log earnings growth $\Delta e_{s,i,t+j}$ capture long-term news about cash flows. $pe_{s,i,t,t+h}$ is a terminal value that captures other long-term influences beyond h years into the future that is not already captured in discount rates and cash flows. The decomposition above holds both ex-ante and ex-post. I consider ex-ante decompositions under subjective expectations $\mathbb{F}_t[\cdot]$:

$$\log q_{s,i,t} = c_q + \mathbb{F}_t[r_{s,i,t,t+h}] - \mathbb{F}_t[e_{s,i,t,t+h}] - \mathbb{F}_t[pe_{s,i,t,t+h}] \quad (\text{A.105})$$

where $\mathbb{F}_t[x_{s,i,t,t+h}]$ denotes the h period ahead subjective expectations of variable x for the firm producing intermediate good i in island s under the time t information set.

D Data Details

This section describes the data sources used in the estimation. I use quarterly data on the macroeconomic and financial series represented in the decomposition from equations (A.37) and (A.38): employment L_t , unemployment U_t , job filling rates q_t , stock returns $r_{t,t+h}$, earnings growth $\Delta e_{t,t+h}$, price-earnings ratio $pe_{t,t+h}$, and earnings-employment ratio $el_{t,t+h}$. For each dependent variable of the decomposition, I also construct their corresponding survey expectations \mathbb{F}_t and machine expectations \mathbb{E}_t .

D.1 Employment

D.1.1 Realized Employment

For realized values of employment, I first construct an annual series for the aggregate number of employees (EMP) of the S&P 500 constituents by using accounting information from the CRSP and Compustat Merged Annual Industrial Files. The data spans 1970 to 2022 and was downloaded from WRDS on July 26, 2023. I interpolate the annual series to a monthly frequency by using the fitted values from real-time regressions of log annual Compustat employment series on the log monthly BLS series for total nonfarm payrolls (PAYEMS). The regressions are estimated over recursively expanding samples from an initial monthly sample that begins on 1970:01 and ends on the month of the data release for each month's total nonfarm payrolls. To ensure that the fitted values do not use future information not available on each data release, I align each monthly BLS nonfarm payroll release with the annual Compustat S&P 500 employment series from the previous calendar year. To obtain a measure of employment L_{t+1} at the beginning of period $t + 1$, I convert the monthly interpolated values to a quarterly frequency by taking the value of the series as of the last month of each calendar quarter. This timing assumption ensures that the measures are consistent with the timing conventions from Section B while still remaining known to firms by the end of period t . Data on nonfarm payrolls was downloaded through FRED on May 15, 2024.

D.1.2 Survey Expectations of Employment

For subjective expectations about employment growth, I use mean point forecasts from the CFO survey, obtained from https://www.richmondfed.org/-/media/RichmondFedOrg/research/national_economy/cfo_survey/current_historical_cfo_data.xlsx. Prior to 2020Q1, the survey asks respondents about their expectations for their company's annual employment growth during the next 12 months: *“Relative to the previous 12 months, what will be your company's PERCENTAGE CHANGE during the next 12 months? (e.g., +3%, 0%, -2%, etc.) [Leave blank if not applicable].”* I interpret the survey to be asking about $\mathbb{F}_t[100 \times (L_{t+1}/L_t - 1)]$, the annual net simple employment growth between the survey month at period t and the same month one year later at period $t + 1$. To obtain survey expectations of annual log employment growth $\mathbb{F}_t[\Delta l_{t+1}]$ from a survey expectation of annual net simple employment growth, I use the approximation

$$\mathbb{F}_t[\Delta l_{t+1}] \approx \log(1 + \mathbb{F}_t[100 \times (L_{t+1}/L_t - 1)]/100) \quad \text{if } t < 2020Q1 \quad (\text{A.106})$$

After 2020Q1, the survey asks respondents about their expectations for their company's employment level for the current and next calendar years: *“What number of employees do you expect to have at the end of calendar years 2021 and 2022? We will assume a value of 0 if left blank.”* Since the reference periods for these forecasts are at the end of the calendar year, I follow De La O and Myers (2021) by interpolating across the end of current and next year expectations to obtain rolling 12-month ahead expectations. For example, if the fiscal year of firm XYZ ends nine months after the survey date, the end of current year expectation has a forecast horizon of 9 months, and the end of next year expectation has a forecast horizon of 21 months. I then obtain the 12-month ahead forecast by interpolating these two forecasts as follows:

$$\mathbb{F}_t[L_{t+1}] = \frac{9}{12}\mathbb{F}_t[L_{t+9/12}] + \frac{3}{12}\mathbb{F}_t[L_{t+21/12}] \quad \text{if } t \geq 2020Q1 \quad (\text{A.107})$$

The survey also asks respondents about the current number of employees at the time of the survey: *“What is your current number of employees? Precise values are preferred. We will assume a value of 0 if left blank.”* I convert the level forecasts to one-year growth forecasts by constructing the log difference of the level forecasts one-year ahead relative to the current number of employees:

$$\mathbb{F}_t[\Delta l_{t+1}] \equiv \mathbb{F}_t[\Delta \log L_{t+1}] \approx \log(\mathbb{F}_t[L_{t+1}]) - \log(L_t) \quad \text{if } t \geq 2020Q1 \quad (\text{A.108})$$

To obtain long-horizon survey expectations of annualized long-horizon log employment growth for $h > 1$ years ahead, I assume that survey respondents expect employment growth to revert back to its historical mean over the next 10 years. Specifically, I interpolate the h year ahead forecast between the 1 year ahead forecast and the historical mean annual log employment growth, with a relative weight on the 1 year ahead forecast that decays linearly over the next 10 years:

$$\mathbb{F}_t[\Delta l_{t+h}] = \frac{10-h}{10-1}\mathbb{F}_t[\Delta l_{t+1}] + \frac{h-1}{10-1}\mathbb{F}_t[\Delta l_{t+10}], \quad h = 1, 2, \dots, 10 \quad (\text{A.109})$$

where $\mathbb{F}_t[\Delta l_{t+10}]$ is the historical mean of the realized log annual employment growth up to the time of the forecast.

The CFO survey panel includes firms that range from small operations to Fortune 500 companies across all major industries. Respondents include chief financial officers, owner-operators, vice presidents and directors of finance, and others with financial decision-making roles. The CFO panel has 1,600 members as of December 2022. I take a stand on the information set of respondents when each forecast was made, and I assume that respondents could have used all data released before they completed the survey. Because the CFO survey releases quarterly forecasts at the end of each quarter, I conservatively set the response deadline for the machine forecast to be the first day of the last month of each quarter (e.g., March 1st).

The data spans the periods 2001Q4 to 2021Q1 and were downloaded on August 8th, 2022. The forecast is not available in 2019Q1, 2019Q4, 2020Q1, and 2020Q2. I impute the missing forecast for 2019Q1 by linearly interpolating between the available forecasts from 2018Q4 and 2019Q2. I impute the missing forecasts for 2019Q4, 2020Q1, and 2020Q2 by interpolating with the nearest available forecast between 2019Q3 and 2020Q3.

D.2 Job Filling Rate

I construct a monthly series for the number of vacancies V_t following Barnichon (2010), by using JOLTS job openings starting 2000:12 (JTS00000000JOL) and extending the series back in time using the help-wanted index before 2000:12. The vacancies data has been downloaded from available on the author's website on May 19, 2024. For realized values of unemployment U_t , I use the BLS monthly series for the unemployment level (UNEMPLOY), downloaded through FRED on May 15, 2024. Labor market tightness $\theta_t = V_t/U_t$ is the ratio between vacancies and unemployment.

I follow Shimer (2012) in constructing the job separation rate δ_t , job finding rate f_t , and job filling rate q_t . Job separation rate is the share of short-term unemployed out of total employment $\delta_t = U_t^s/L_t$, where U_t^s is the BLS series for the number of unemployed less than 5 weeks (UEMPLT5) that was downloaded through FRED on May 15, 2024. The job finding rate is:

$$f_t = 1 - \frac{U_t - U_t^s}{U_{t-1}} \quad (\text{A.110})$$

The expression for the job finding rate follows from the unemployment accumulation equation:

$$U_t = (1 - f_t)U_{t-1} + U_t^s \quad (\text{A.111})$$

which states that unemployment U_t consists of either the previously unemployed U_{t-1} who did not find a job $(1 - f_t)$, or the short-term unemployed U_t^s that lost a job during the current period. The job filling rate is defined as the share of filled vacancies $f_t V_t$ out of unemployment U_t :

$$q_t = \frac{f_t}{\theta_t} = \frac{f_t V_t}{U_t} \quad (\text{A.112})$$

I first construct the job filling rate q_t at the monthly frequency. To remove high-frequency fluctuations that likely reflect measurement errors, I time-aggregate the monthly series to a quarterly frequency by taking a 3-month trailing average that ends on the first month of each calendar quarter. This timing assumption ensures that the survey and machine expectations in the variance decomposition do not use advance information about job filling rates that were not published at the time of each forecast. To ensure that all variables used in the variance decomposition are stationary, I follow Shimer (2012) by detrending the quarterly job filling rate q_t using an HP filter with a smoothing parameter of 10^5 .

D.3 Wages

To assess the cyclical nature of subjective wage expectations, I use publicly available survey and macroeconomic data to construct measures of actual real wage growth, subjective wage expectations, and unemployment rate changes. The Livingston Survey (semi-annual, 1961S1–2022S2), the CFO Survey (quarterly, 2001Q4–2022Q4), and the Survey of Consumer Expectations (SCE) (monthly, 2015M5–2022M12) provide the necessary data.

I derive subjective wage growth expectations from median consensus forecasts of nominal wage growth in these surveys. The Livingston Survey forecasts are deflated using its own median CPI inflation forecast, while the CFO and SCE survey forecasts are deflated using CPI inflation expectations from the Survey of Professional Forecasters (SPF).

To account for the possibility that wages depend on the economic conditions at the start of the job, I use survey expectations from the SCE to measure the user cost of labor UC_t^W under subjective expectations. In the search and matching model, the user cost of labor is the difference in the present value of wages between two firm-worker matches that are formed in two consecutive periods. Existing work measures the user cost of labor under full information rational expectations and finds that the user cost is more cyclical than the flow wage, suggesting that workers hired in recessions earn lower wages not only when hired but also in subsequent periods (Kudlyak, 2014; Bils et al., 2023). The survey-based measure in this paper relaxes the rational expectations assumption maintained in existing work. Consider the free-entry condition in the search and matching model:

$$\frac{\kappa}{q_t} = J_{t,t} \tag{A.113}$$

where a firm must pay a per vacancy cost of κ and vacancies are filled with probability q_t . $J_{t,\tau}$ is the value of a firm with a worker at time τ such that the productive match started at time t :

$$J_{t,\tau} \equiv z_t - w_{t,t} + \sum_{\tau=t+1}^{\infty} (\beta(1 - \delta))^{\tau-t} \mathbb{F}_t[z_\tau - w_{t,\tau}] \tag{A.114}$$

where $\mathbb{F}_t[\cdot]$ denotes subjective expectations based on survey data. $\beta = 0.9569$ is a discount factor and $\delta = 0.295$ is the probability that an employment relationship is terminated, both from Kudlyak (2014). Each period τ , a firm-worker match produces a per period output of z_τ and

an employed worker received wage $w_{t,\tau}$ where t is the period when the worker is hired. $w_{t,t}$ is the new-hire wage. Note that the free entry condition is only required to hold for newly created matches for $\tau = t$. The expected difference between the firm's value of a newly created match in time t and the discounted value of a newly created match in period $t + 1$ is

$$J_{t,t} - \beta(1 - \delta)\mathbb{E}_t[J_{t+1,t+1}] = z_t - w_{t,t} + \sum_{\tau=t+1}^{\infty} (\beta(1 - \delta))^{\tau-t}\mathbb{E}_t[z_{\tau} - w_{t,\tau}] \quad (\text{A.115})$$

$$- \beta(1 - \delta)\mathbb{E}_t \left[z_{t+1} - w_{t+1,t+1} + \sum_{\tau=t+2}^{\infty} (\beta(1 - \delta))^{\tau-(t+1)}\mathbb{E}_{t+1}[z_{\tau} - w_{t+1,\tau}] \right] \quad (\text{A.116})$$

Apply the Law of Iterated Expectations and collect terms

$$J_{t,t} - \beta(1 - \delta)\mathbb{E}_t[J_{t+1,t+1}] = z_t - w_{t,t} - \sum_{\tau=t+1}^{\infty} (\beta(1 - \delta))^{\tau-t}\mathbb{E}_t[w_{t,\tau} - w_{t+1,\tau}] \quad (\text{A.117})$$

Substitute the free-entry condition to the left-hand side

$$\underbrace{\frac{\kappa}{q_t} - \beta(1 - \delta)\mathbb{E}_t \left[\frac{\kappa}{q_{t+1}} \right]}_{\text{Non-wage component of user cost } UC_t^V} = \underbrace{z_t}_{\text{Benefit}} - \underbrace{\left[w_{t,t} + \sum_{\tau=t+1}^{\infty} (\beta(1 - \delta))^{\tau-t}\mathbb{E}_t[w_{t,\tau} - w_{t+1,\tau}] \right]}_{\text{Wage component of user cost } UC_t^W} \quad (\text{A.118})$$

The equation shows that the firm faces two sources of costs from a match: wage payments to a worker UC_t^W and vacancy opening costs UC_t^V . The firm creates jobs as long as the marginal benefit from adding a worker exceeds the user cost of labor. It is worth noting that the wage component of the user cost of labor UC_t^W , not the wage $w_{t,t}$, is the allocative price of labor.

I use worker-level data from the Survey of Consumer Expectations (SCE) to construct the user cost of labor UC_t^W under the survey respondents' subjective expectations. The SCE asks respondents about: the *month and year on which their current employment relationship started* (i.e., t in $w_{t,\tau}$); *“annual earnings, before taxes and other deductions, on your [current/main] job”* ($w_{t,\tau}$); short-term expectations on what their *“annual earnings will be in 4 months”* ($\mathbb{E}_t[w_{t,t+\frac{4}{12}}]$) and long-term expectations on *“annual earnings to be at your current job in 10 years”* ($\mathbb{E}_t[w_{t,t+10}]$). I obtain survey expectations about medium-term earnings between 4 months to 10 years by linearly interpolating between the two horizons:

$$\mathbb{E}_t[w_{t,t+h}] = \frac{10 - h}{10 - \frac{4}{12}}\mathbb{E}_t[w_{t,t+\frac{4}{12}}] + \frac{h - \frac{4}{12}}{10 - \frac{4}{12}}\mathbb{E}_t[w_{t,t+10}], \quad h = 1, 2, \dots, 10 \quad (\text{A.119})$$

The user cost of labor formulation assumes infinitely lived firms and workers, while empirical data are inherently finite. I truncate the horizon at 10 years given the availability of the survey data. Longer horizons reduce the weight of future terms due to discounting and job separations. In addition, if unemployment follows a mean-reverting process, wages in long-term employment relationships will eventually converge to the long-term mean, which after discounting would limit the size of very long-term influences (Kudlyak, 2014).

I measure actual real wage growth using two BLS wage series. The Livingston Survey forecasts target annual log real wage growth based on average weekly earnings of production and nonsupervisory employees in manufacturing (CES30000000030). The CFO and SCE surveys target annual log real wage growth based on average hourly earnings of private-sector employees (CEU0500000008). I deflate nominal wages using the Consumer Price Index (CPIAUCSL) to adjust for purchasing power.

For unemployment rates used to assess the cyclicalities of wages, I use both actual data and subjective forecasts. The actual seasonally adjusted U.S. unemployment rate (UNRATE) comes from the BLS Current Population Survey (CPS), while subjective unemployment expectations are derived from median consensus SPF forecasts of future unemployment rates.

D.4 Stock Returns

D.4.1 Realized Stock Returns

To measure stock market returns, I use monthly data on CRSP value-weighted returns including dividends (VWRETD) from the Center for Research in Security Prices (CRSP). I compute annualized log stock returns by compounding the monthly returns using $r_{t+h} \equiv (1/h) \sum_{j=1}^{12h} \log(1 + VWRETD_{t+j/12})$. The data was downloaded from WRDS on February 12, 2023. When evaluating the MSE ratios of the machine relative to that of a benchmark survey, I compute machine forecasts for either annual CRSP returns or S&P 500 price growth depending on which value most closely aligns with the concept that survey respondents are asked to predict. To measure one-year stock market price growth, I use the one-year log cumulative growth rate of the S&P 500 index, $\Delta p_{t+1} \equiv \log(P_{t+1}/P_t)$. The monthly S&P index series spans the period 1957:03 to 2022:12 and was downloaded from WRDS on January 24, 2024 from the Annual Update data of the Index File on the S&P 500.

D.4.2 Survey Expectations of Stock Returns

CFO Survey I use survey forecasts of S&P 500 stock returns from the CFO survey to measure subjective return expectations. The CFO survey is a quarterly survey that asks respondents about their expectations for the S&P 500 return over the next 12 months and 10 years ahead, obtained from https://www.richmondfed.org/-/media/RichmondFedOrg/research/national_economy/cfo_survey/current_historical_cfo_data.xlsx. I use the mean point forecast for the value of the “most likely” future stock return in the estimation. More specifically, the survey asks the respondent “*over the next 12 months, I expect the average annual S&P 500 return will be: Most Likely: I expect the return to be: ___%*”. The survey question for stock return expectations 10 years ahead is “*over the next 10 years, I expect the average annual S&P 500 return will be: Most Likely: I expect the return to be: ___%*”.

The CFO survey panel includes firms that range from small operations to Fortune 500 companies across all major industries. Respondents include chief financial officers, owner-operators,

vice presidents and directors of finance, and others with financial decision-making roles. The CFO panel has 1,600 members as of December 2022.

I take a stand on the information set of respondents when each forecast was made, and I assume that respondents could have used all data released before they completed the survey. Because the CFO survey releases quarterly forecasts at the end of each quarter, I conservatively set the response deadline for the machine forecast to be the first day of the last month of each quarter (e.g., March 1st).

The data spans the periods 2001Q4 to 2021Q1 and were downloaded on August 8th, 2022. Mean point forecasts before 2020Q3 are available in column `sp_1_exp` of sheet `through_Q1_2020`; mean point forecasts from 2020Q3 and onwards are available in column `sp_12moexp_2` of sheet `CFO_SP500`. The forecast is not available in 2019Q1, 2019Q4, 2020Q1, and 2020Q2. I impute the missing forecast for 2019Q1 by linearly interpolating between the available forecasts from 2018Q4 and 2019Q2. I impute the missing forecasts for 2019Q4, 2020Q1, and 2020Q2 by interpolating with the nearest available forecast between 2019Q3 and 2020Q3.

Following Nagel and Xu (2022), I assume that the forecasted S&P 500 return includes dividends and capture expectations about annualized cumulative simple net returns compounded from time t to $t+h$, i.e., $\mathbb{F}_t[R_{t,t+h}]$. To obtain survey expectations of log returns $\mathbb{F}_t[\log(1+r_{t,t+h})]$ from a survey expectation of net simple returns $\mathbb{F}_t[R_{t,t+h}]$, I use the approximation $\mathbb{F}_t[\log(1+r_{t,t+h})] \approx \log(1+\mathbb{F}_t[R_{t,t+h}])$.

To obtain long-horizon survey expectations of annualized cumulative log S&P 500 returns over the next $1 < h < 10$ years, I interpolate the forecasts across annualized 1 year and 10 year cumulative log return expectations:

$$\mathbb{F}_t[r_{t,t+h}] = \frac{10-h}{10-1}\mathbb{F}_t[r_{t,t+1}] + \frac{h-1}{10-1}\mathbb{F}_t[r_{t,t+10}], \quad h = 1, 2, \dots, 10 \quad (\text{A.120})$$

Finally, I use the difference between the cumulative annualized long-horizon log return expectations between adjacent years (i.e., $\mathbb{F}_t[r_{t,t+h-1}]$ and $\mathbb{F}_t[r_{t,t+h}]$) to obtain $\mathbb{F}_t[r_{t+h}]$, the time t survey expectation of forward one-year log stock returns h years ahead:

$$\mathbb{F}_t[r_{t+h}] = h \times \mathbb{F}_t[r_{t,t+h}] - (h-1) \times \mathbb{F}_t[r_{t,t+h-1}], \quad h = 1, 2, \dots, 10 \quad (\text{A.121})$$

Gallup/UBS Survey The UBS/Gallup is a monthly survey of one-year-ahead stock market return expectations. I use the mean point forecast in our estimation and compare these to machine forecasts of the annual CRSP return. Gallup conducted 1,000 interviews of investors during the first two weeks of every month and results were reported on the last Monday of the month. The first survey was conducted on 1998:05. Until 1992:02, the survey was conducted quarterly on 1998:05, 1998:09, and 1998:11. The data on 1998:06, 1998:07, 1998:08, 1998:10, 1998:12, 1999:01, and 2006:01 are missing because the survey was not conducted on these months. I follow Adam et al. (2021) in starting the sample after 1999:02 due to missing values at the beginning of the sample.

For each month when the survey was conducted, respondents are asked about the return they expect on their own portfolio. The survey question is “*What overall rate of return do you expect to get on your portfolio in the next twelve months?*” Before 2003:05, respondents are also asked about the return they expect from an investment in the stock market during the next 12 months. The survey question is “*Thinking about the stock market more generally, what overall rate of return do you think the stock market will provide investors during the coming twelve months?*” For each month, I calculate the average expectations of returns on their own portfolio and returns on the market index. When calculating the average, survey respondents are weighted by the weight factor provided in the survey. I exclude extreme observations where a respondent reported expected returns higher than 95% or lower than -95% on either their own portfolio or the market index.

In order to construct a consistent measure of stock market return expectations over the entire sample period, I impute missing market return expectations using the fitted values from two regressions. First, I impute missing values during 1999:02-2005:12 and 2006:02-2007:10 with the fitted value from regressing expected market returns on own portfolio expectations contemporaneously, where the regression is estimated using the part of the sample where both are available. Second, I impute the one missing observation in both market and own portfolio return expectations for 2006:01 with the fitted value from regressing the market return expectations on the lagged own portfolio return expectations, where the coefficients are estimated using part of the sample where both are available, and the fitted value combines the estimated coefficients with lagged own portfolio expectations data from 2005:12.

Following Nagel and Xu (2022), I assume that the forecasted stock market return includes dividends and capture expectations about annual simple net stock returns $\mathbb{F}_t[R_{t+1}]$. To obtain survey expectations of annual log returns $\mathbb{F}_t[\log(1 + r_{t+1})]$ from a survey expectation of annual net simple returns $\mathbb{F}_t[R_{t+1}]$, I use the approximation $\mathbb{F}_t[\log(1 + r_{t+1})] \approx \log(1 + \mathbb{F}_t[R_{t+1}])$. After applying all the procedures, the Gallup market return expectations series spans the periods 1999:02 to 2007:10. The data were downloaded on August 1st, 2024 from Roper iPoll: <http://ropercenter.cornell.edu/ubs-index-investor-optimism/>.

I take a stand on the information set of respondents when each forecast was made, and I assume that respondents could have used all data released before they completed the survey. Since interviews are in the first two weeks of a month (e.g., February), I conservatively set the response deadline for the machine forecast to be the first day of the survey month (e.g., February 1st), implying that I allow the machine to use information only up through the end of the previous month (e.g., through January 31st). This ensures that the machine only sees information that would have been available to all UBS/Gallup respondents for that survey month (February). This approach is conservative in the sense that it handicaps the machine, since all survey respondents who are being interviewed during the next month would have access to more timely information than the machine. Since the survey asks about the “one-year-ahead” I interpret the question to be asking about the forecast period spanning from the current survey month to the same month

one year ahead.

Michigan Survey of Consumers (SOC) The SOC contains approximately 50 core questions, and a minimum of 500 interviews are conducted by telephone over the course of the entire month, each month. Table 20 of the SOC reports the probability of an increase in stock market in next year. The survey question was *“The next question is about investing in the stock market. Please think about the type of mutual fund known as a diversified stock fund. This type of mutual fund holds stock in many different companies engaged in a wide variety of business activities. Suppose that tomorrow someone were to invest one thousand dollars in such a mutual fund. Please think about how much money this investment would be worth one year from now. What do you think the percent chance that this one thousand dollar investment will increase in value in the year ahead, so that it is worth more than one thousand dollars one year from now?”*

When I use this survey forecast to compare to machine forecasts, I impute a point forecast for stock market returns using the method described in Section D.4.2 below. I compare the imputed point forecast to machine forecasts of CRSP returns.

For the SOC, interviews are conducted monthly typically over the course of an entire month. (In rare cases, interviews may commence at the end of the previous month, as in February 2018 when interviews began on January 31st 2018.) I take a stand on the information set of respondents when each forecast was made, and I assume that respondents could have used all data released before they completed the survey. Since interviews are almost always conducted over the course of an entire month (e.g., February), I conservatively set the response deadline for the machine forecast to be the first day of the survey month (e.g., February 1st), implying that I allow the machine to use information only up through the end of the previous month (e.g., through January 31st). This ensures that the machine only sees information that would have been available to all respondents for that survey month (February). This approach is conservative in the sense that it handicaps the machine, since all survey respondents who are being interviewed during the next month would have access to more timely information than the machine. Since the survey asks about the “year ahead” I interpret the question to be asking about the forecast period spanning the period running from the current survey month to the same month one year ahead. The data spans 2002:06 to 2021:12. The SOC responses were obtained from <https://data.sca.isr.umich.edu/data-archive/mine.php> and downloaded on August 13th, 2022.

Livingston Survey Stock Price Forecast I obtain the Livingston Survey S&P 500 index forecast (SPIF) from the Federal Reserve Bank of Philadelphia, URL: <https://www.philadelphiafed.org/surveys-and-data/real-time-data-research/livingston-historical-data>, and use the mean values in our structural and forecasting models. I compare the one-year growth in these forecasts to machine forecasts of S&P 500 price growth. Our sample spans 1947:06 to 2021:06. The forecast series were downloaded on September 20, 2021.

The survey provides semi-annual forecasts on the level of the S&P 500 index. Participants are asked to provide forecasts for the level of the S&P 500 index for the end of the current survey month, 6 months ahead, and 12 months ahead. I use the mean of the respondents' forecasts each period, where the sample is based on about 50 observations. Most of the survey participants are professional forecasters with “formal and advanced training in economic theory and forecasting and use econometric models to generate their forecasts.”

Participants receive questionnaires for the survey in May and November, after the Consumer Price Index (CPI) data release for the previous month. All forecasts are typically submitted by the end of the respective month of May and November. The results of the survey are released near the end of the following month, on June and December of each calendar year. The exact release dates are available on the Philadelphia Fed website, at the header of each news release. I take a stand on the information set of the respondents when each forecast was made by assuming that respondents could have used all data released before they completed the survey. Since all forecasts are typically submitted by the end of May and November of each calendar year, I set the response deadline for the machine forecast to be the first day of the last month of June and December, implying that I allow the machine to use information only up through the end of the May and November.

I follow Nagel and Xu (2021) in constructing one-year stock price growth expectations from the level forecasts. Starting from June 1992, I use the ratio between the 12-month level forecast (SPIF_12M_t) and 0-month level nowcasts (SPIF_ZM_t) of the S&P 500 index. Before June 1992, the 0-month nowcast is not available. Therefore I use the annualized ratio between the 12-month (spi12_t) and 6-month (spi6_t) level forecast of the S&P 500 index

$$\mathbb{F}_t^{(Liv)} \left[\frac{P_{t+1}}{P_t} \right] \approx \begin{cases} \frac{\mathbb{F}_t^{(Liv)}[P_{t+1}]}{\mathbb{F}_t^{(Liv)}[P_t]} = \frac{\text{SPIF_12M}_t}{\text{SPIF_ZM}_t} & \text{if } t \geq 1992M6 \\ \left(\frac{\mathbb{F}_t^{(Liv)}[P_{t+1}]}{\mathbb{F}_t^{(Liv)}[P_{t+6}]} \right)^2 = \left(\frac{\text{spi12}_t}{\text{spi6}_t} \right)^2 & \text{if } t < 1992M6 \end{cases} \quad (\text{A.122})$$

where P_t is the S&P 500 index and t indexes the survey's response deadline. To obtain a survey expectation of the log change in price growth I use the approximation:

$$\mathbb{F}_t(\Delta p_{t+1}) \approx \log(\mathbb{F}_t[P_{t+1}]) - \log(P_t)$$

Conference Board (CB) Survey Respondents provide the categorical belief of whether they expect stock prices to “increase,” “decrease,” or stay the “same” over the next year. Since the survey asks respondents about stock prices in the “year ahead,” I interpret the question to be asking about the forecast period from the end of the current survey month to the end of the same month one year ahead. When we use this qualitative survey forecast to compare to machine forecasts, we impute a point forecast for stock market returns using the method described in Section D.4.2 below. I compare the imputed point forecast to machine forecasts of CRSP returns.

The survey is conducted monthly and I use the survey responses over 1987:04 to 2022:08. The data was downloaded on September 26, 2022. The survey uses an address-based mail sample

design. Questionnaires are mailed to households on or about the first of each month. Survey responses flow in throughout the collection period, with the sample close-out for preliminary estimates occurring around the 18th of the month. Any responses received after then are used to produce final estimates for the month, which are published with the following month’s data. Conversations with those knowledgeable about the survey suggested that most panelists respond early. Any responses received after around the 20th of the month—regardless of when they are filled out—are included in the final (but not preliminary) numbers.

I take a stand on the information set of the respondents when each forecast was made by assuming that respondents could have used all data released before they completed the survey. Since questionnaires reach households on or about the first of each month (e.g., February 1st) and most respondents respond early, I conservatively set the response deadline for the machine forecast to be the first day of the survey month (e.g., February 1st), implying that I allow the machine to use information only up through the end of the previous month (e.g., January 31st).

Converting Qualitative Forecasts to Point Forecasts (SOC and CB) I use the SOC probability to impute a quantitative point forecast of stock returns using a linear regression of CFO point forecasts for returns onto the SOC probability of a price increase. The SOC asks respondents about the percent chance that an investment will “increase in value in the year ahead.” I interpret this as asking about the ex dividend value, i.e., about price price growth. The CFO survey is conducted quarterly, where the survey quarters span 2001Q4 to 2021Q1. The SOC survey is conducted monthly, where survey months span 2002:06 to 2021:12. Since the CFO is a quarterly survey, the regression is estimated in real-time over a quarterly overlapping sample. Since the CFO survey is conducted during the last month of the quarter while the SOC is conducted monthly, I align the survey months between CFO and SOC by regressing the quarterly CFO survey point forecast with the qualitative SOC survey response during the last month of the quarter.

Since the SOC survey question is interpreted as asking about S&P 500 price growth while the CFO survey question asks about stock returns including dividends, I follow Nagel and Xu (2021) in subtracting the current dividend yield of the CRSP value weighted index from the CFO variable before running the regression. After estimating the regression, I then add back the dividend yield to the fitted value to obtain an imputed SOC point forecast of stock returns including dividends.

Specifically, at time t , I assume that the CFO forecast of stock returns, $\mathbb{F}_t^{\text{CFO}}[r_{t,t+1}]$, minus the current dividend yield, D_t/P_t , is related to the contemporaneous SOC probability of an increase in the stock market next year, $P_{t,t+1}^{\text{SOC}}$, by:

$$\mathbb{F}_t^{\text{CFO}}[r_{t,t+1}] - D_t/P_t = \beta_0 + \beta_1 P_{t,t+1}^{\text{SOC}} + \epsilon_t.$$

The final imputed SOC point forecast is constructed as $\mathbb{F}_t^{\text{SOC}}[r_{t,t+1}] = \hat{\beta}_0 + \hat{\beta}_1 P_{t,t+1}^{\text{SOC}} + D_t/P_t$. I first estimate the coefficients of the above regression over an initial overlapping sample of 2002Q2 to

2004Q4, where the quarterly observations from the CFO survey is regressed on the SOC survey responses from the last month of each calendar quarter. Using the estimated coefficients and the SOC probability from 2005:03 gives us the point forecast of the one-year stock return from 2005Q1 to 2006Q1. I then re-estimate this equation, recursively, adding one quarterly observation to the end of the sample at a time, and storing the fitted values. This results in a time series of SOC point forecasts $\mathbb{F}_t^{\text{SOC}}[r_{t,t+1}]$ spanning 2005Q1 to 2021Q1.

The same procedure is done for the Conference Board Survey, except I replace $P_{t,t+1}^{\text{SOC}}$ by $P_{t,t+1}^{\text{CB}}$, a ratio of the proportion of those who respond with “increase” to the sum of “decrease” and “same.” The CB survey asks respondents to provide the categorical belief of whether they expect stock prices to “increase,” “decrease,” or stay the “same” over the next year. I interpret this as asking about price price growth. Since the CB survey question is interpreted as asking about S&P 500 price growth while the CFO survey question asks about stock returns including dividends, I follow Nagel and Xu (2021) in subtracting the current dividend yield of the CRSP value weighted index from the CFO variable before running the regression. After estimating the regression, I then add back the dividend yield to the fitted value to obtain an imputed CB point forecast of stock returns including dividends.

The CFO survey is conducted quarterly, where the survey quarters span 2001Q4 to 2021Q1. The CB survey is conducted monthly, where survey months span 1987:04 to 2022:08. The regression is first estimated over an initial overlapping sample of 2001Q4 to 2004Q4, where the quarterly observations from the CFO survey is regressed on the CB survey responses from the last month of each calendar quarter. Using the estimated coefficients and the CB survey response $P_{t,t+1}^{\text{CB}}$ from 2005:03 gives us the point forecast of the stock return from 2005Q1 to 2006Q1. I then re-estimate this equation, recursively, adding one observation to the end of the sample at a time, and storing the fitted values. This results in a time series of CB point forecasts $\mathbb{F}_t^{\text{CB}}[r_{t,t+1}]$ over 2005Q1 to 2021Q1.

Nagel and Xu Individual Investor Expectations Nagel and Xu (2021)’s individual investor expectations series for returns covers 1972-1977 (Annual) and 1987Q2-2021Q4 (Quarterly) and combine data from the following surveys:

1. UBS/Gallup: 1998:06-2007:10; Survey captures respondents’ expected stock market returns, in percent, over a 1-year horizon.
2. Michigan Survey of Consumers (SOC): 2002:04-2022:12; Respondents provide the probability of a rise in the stock market over a 1-year horizon.
3. Conference Board (CB): 1987:04-2022:08; Respondents provide the categorical opinion whether they expect stock prices to rise, or stay about where they are, or decline over the next year.
4. Vanguard Research Initiative (VRI): 2014:08; Survey captures respondents’ expected stock market returns, in percent, over a 1-year horizon.

5. Roper: 1974-1977, annual, observed June of each calendar year; Respondents provide the categorical opinion whether they expect stock prices to rise, or stay about where they are, or decline over the next year.
6. Lease, Lewellen, and Schlarbaum (1974, 1977): 1972-1973, annual, observed July of each calendar year; Survey captures respondents' expected stock market returns, in percent, over a 1-year horizon.

NX arrive at their final series through the following imputation steps. UBS/Gallup measures the investors' expected stock market returns most closely, but it covers a relatively short period. SOC and CB cover a relatively longer period, but they are qualitative forecasts that need to be converted to point forecasts.

1. Start with UBS/Gallup for 1998:06-2007:10 and VRI for 2014:08 since they capture the respondents' expected stock returns relatively closely (other surveys only provide qualitative measures).
2. Regress SOC on UBS/Gallup and VRI using periods of overlapping coverage (2002:04-2007:10). Use the fitted values from this regression to impute missing data for 2007:11-2022:12 (excluding 2014:08).
3. Regress CB on UBS/Gallup and VRI using periods of overlapping coverage (1998:06-2007:10). Use the fitted values from this regression to impute missing data for 1987:04-1998:05 (using CB) and 1974-1977 (using Roper).
4. Use the coefficients from regressing CB on UBS/Gallup and VRI (from step 3) to compute fitted values that convert the probabilistic forecast from Roper into point forecasts of stock returns.
5. Convert expected returns to expected excess returns by subtracting the average 1-year Treasury yield measured at the beginning of the survey month.
6. Aggregate monthly series to a quarterly frequency by taking the average expectation within calendar quarters.

D.5 Risk-Free Rates

Realized Risk-Free Rates As a measure of realized risk-free rates r_t^f , I obtain daily series for the annualized three-month Treasury bill rate (DTB3), downloaded from FRED on May 15, 2024. To match the definition used as the target variable in the Survey of Professional Forecasters (SPF), I time-aggregate the daily realized risk-free rate series to a quarterly frequency by taking the quarterly average, as discussed below.

Survey Expectations of Risk-Free Rates I obtain subjective expectations about risk-free rates from median forecasts for the annualized three-month Treasury bill rate from the Survey of Professional Forecasters (SPF). The SPF provides forecasts at the one and ten year horizons.

For one year ahead forecasts (TBILL), respondents are asked to provide quarterly forecasts of the quarterly average three-month Treasury bill rate, in percentage points, where the forecasts are for the quarterly average of the underlying daily levels. I interpret the survey to be asking about annual net simple rates $\mathbb{F}_t[R_{t,t+1}^f]$, and approximate the expected log risk-free rate as $\mathbb{F}_t[r_{t,t+1}^f] \approx \log(1 + \mathbb{F}_t[R_{t,t+1}^f])$.

For ten year ahead forecasts (BILL10), respondents are asked to provide forecasts for the annual-average rate of return to three-month Treasury bills over the next 10 years, in percentage points. The ten year ahead forecasts are available only for surveys conducted in the first quarter of each calendar year. I interpret the survey to be asking about annualized cumulative net simple rates compounded from the survey quarter to the same quarter that is ten years after the survey year $\mathbb{F}_t[R_{t,t+10}^f]$, and approximate the expected log risk-free rate as $\mathbb{F}_t[r_{t,t+10}^f] \approx \log(1 + \mathbb{F}_t[R_{t,t+10}^f])$.

To obtain long-horizon survey expectations of annualized log three-month Treasury bill rates over the next $1 < h < 10$ years, I interpolate the forecasts across annualized 1 year and 10 year return expectations:

$$\mathbb{F}_t[r_{t,t+h}^f] = \frac{10-h}{10-1} \mathbb{F}_t[r_{t,t+1}^f] + \frac{h-1}{10-1} \mathbb{F}_t[r_{t,t+10}^f], \quad h = 1, 2, \dots, 10 \quad (\text{A.123})$$

Finally, I use the difference between the cumulative annualized long-horizon log three-month Treasury bill rate expectations between adjacent years (i.e., $\mathbb{F}_t[r_{t,t+h-1}^f]$ and $\mathbb{F}_t[r_{t,t+h}^f]$) to obtain $\mathbb{F}_t[r_{t+h}^f]$, the time t survey expectation of annualized forward log three-month Treasury bill rate h years ahead:

$$\mathbb{F}_t[r_{t+h}^f] = h \times \mathbb{F}_t[r_{t,t+h}^f] - (h-1) \times \mathbb{F}_t[r_{t,t+h-1}^f], \quad h = 1, 2, \dots, 10 \quad (\text{A.124})$$

The surveys are sent out at the end of the first month of each quarter, and they are collected in the second or third week of the middle month of each quarter. When constructing machine learning forecasts for the risk-free rate, I assume that forecasters could have used all data released before the survey deadlines for the SPF, which are posted online at the Federal Reserve Bank of Philadelphia website. Since surveys are typically sent out at the end of the first month of each quarter, I make the conservative assumption that respondents only had data released by the first day of the second month of each quarter.

D.6 Earnings

D.6.1 Realized Earnings

To measure corporate earnings, I use quarterly S&P 500 earnings per share (EPS) data from S&P Global <https://www.spglobal.com/spdji/en/documents/additional-material/sp-500-eps-est.xlsx>. The EPS series starts from 1988Q2. To extend the sample backward, I use monthly EPS data on the

S&P 500 from Robert Shiller's data depository at URL: http://www.econ.yale.edu/~shiller/data/ie_data.xls. These are monthly EPS data equal to the sum over the trailing 12 months, computed from the S&P four-quarter trailing totals. I use data for this series from 1959:01 to 1988:03. (There are no quarterly EPS observations recorded publically prior to 1988:03.) To obtain a single quarterly earnings series extending backward to 1959Q1, I employ a recursive process that combines these two series. Specifically, let time t be measured in months. Shiller's series provides a monthly series of earnings over the past 12 months:

$$em_t = e_t + e_{t-1} + \dots + e_{t-11} \quad (\text{A.125})$$

Starting from 1988Q2, I also have a quarterly series of earnings over the quarter:

$$eq_t = e_t + e_{t-1} + e_{t-2} \quad (\text{A.126})$$

if $t = \{3, 6, 9, 12\}$, $eq_t = 0$ otherwise. Suppose that I am interested in earnings for 1988Q1:

$$em_{1988M12} = \underbrace{e_{1988M12} + e_{1988M11} + e_{1988M10} + \dots + e_{1988M3} + e_{1988M2} + e_{1988M1}}_{eq_{1988M12}} \quad (\text{A.127})$$

$$= eq_{1988M12} + eq_{1988M9} + eq_{1986M6} + eq_{1988M3} \quad (\text{A.128})$$

We can then compute implied earnings for 1988Q1 as

$$eq_{1988M3} = em_{1988M12} - [eq_{1988M12} + eq_{1988M9} + eq_{1988M6}]. \quad (\text{A.129})$$

We can then use the same formula recursively to obtain earnings before 1988, i.e., with

$$em_{1988M9} = \underbrace{e_{1988M9} + e_{1988M8} + e_{1988M7} + \dots + e_{1987M12} + e_{1987M11} + e_{1987M10}}_{eq_{1988M9}} \quad (\text{A.130})$$

which gives the implied earnings for 1987Q4. I continue recursively, working backward to the beginning of the sample in 1959Q1. In the S&P Global dataset there is one observation in 2008Q4 with a negative EPS. Since I need to ultimately compute earnings growth rates, I remove this single negative observation by replacing the 2008Q4 EPS observation with the Shiller 12-month trailing total EPS observation for 2008Q4. To convert EPS to total earnings, I next multiply the resulting quarterly EPS series by the quarterly S&P 500 divisor available at URL: https://ycharts.com/indicators/sp_500_divisor. Finally, to obtain a monthly S&P 500 earnings series, I linearly interpolate the resulting quarterly total earnings series. The final monthly total earnings series spans the period 1959:03 to 2021:12. The EPS data from S&P Global, Shiller, and the divisor data were downloaded on March 13, 2022.

D.6.2 Survey Expectations of Earnings

I obtain monthly survey data for the median analyst earnings per share forecast and actual earnings per share from the Institutional Brokers Estimate System (IBES) via Wharton Research Data Services (WRDS). The data spans the period 1976:01 to 2021:12 and was downloaded on October 2022.

Short-Term Growth (STG) Expectations I build measures of aggregate S&P 500 earnings expectations growth using the constituents of the S&P 500 at each point in time following De La O and Myers (2021). I first construct expected earnings expectations for aggregate earnings h -months-ahead as:

$$\mathbb{F}_t[E_{t+h}] = \Omega_t \left[\sum_{i \in x_{t+h}} \mathbb{F}_t[EPS_{i,t+h}] S_{i,t} \right] / Divisor_t, \quad (\text{A.131})$$

where \mathbb{F} is the median analyst survey forecast, E is aggregate S&P 500 earnings, EPS_i is earning per share of firm i among all S&P 500 firms x_{t+h} for which I have forecasts in IBES for $t+h$, S_i is shares outstanding of firm i , and $Divisor_t$ is calculated as the S&P 500 market capitalization divided by the S&P 500 index. I obtain the number of outstanding shares for all companies in the S&P500 from Compustat. All data from Compustat were downloaded on November 17th, 2022. IBES estimates are available for most but not all S&P 500 companies. Following De La O and Myers (2021), I multiply this aggregate by Ω_{t+h} , a ratio of total S&P 500 market value to the market value of the forecasted companies at $t+h$ to account for the fact that IBES does not provide earnings forecasts for all firms in the S&P 500 in every period.

IBES database contains earning forecasts up to five annual fiscal periods (FY1 to FY5) and as a result, I interpolate across the different horizons to obtain the expectation over the next 12 months. This procedure has been used in the literature, including De La O and Myers (2021). Specifically, if the fiscal year of firm XYZ ends nine months after the survey date, I have a 9-month earning forecast $\mathbb{F}_t[E_{t+9}]$ from FY1 and a 21-month forecast $\mathbb{F}_t[E_{t+21}]$ from FY2. I then obtain the 12-month ahead forecast by interpolating these two forecasts as follows,

$$\mathbb{F}_t[E_{t+12}] = \frac{9}{12} \mathbb{F}_t[E_{t+9}] + \frac{3}{12} \mathbb{F}_t[E_{t+21}]. \quad (\text{A.132})$$

For the forecasting performance estimates, I use quarterly data. To convert the monthly forecast to quarterly frequency, I use the forecast made in the middle month of each quarter, and construct one-year earnings expectations from 1976Q1 to 2021Q4 and the earning expectation growth is calculated as an approximation following following De La O and Myers (2021):

$$\mathbb{F}_t(\Delta e_{t+12}) \approx \ln(\mathbb{F}_t[E_{t+12}]) - e_t \quad (\text{A.133})$$

where e_t is log earnings for S&P 500 at time t calculated as $e_t = \log(EPS_t \cdot Divisor_t)$, where EPS_t is the earnings per share for the S&P 500 obtained from Shiller's data depository and S&P Global, as described above.

Long-Term Growth (LTG) Expectations I construct long term expected earnings growth (LTG) for the S&P 500 following Bordalo et al. (2019). Specifically, I obtain the median firm-level LTG forecast from IBES, and aggregate the value-weighted firm-level forecasts,

$$LTG_t = \sum_{i=1}^S LTG_{i,t} \frac{P_{i,t} Q_{i,t}}{\sum_{i=1}^S P_{i,t} Q_{i,t}} \quad (\text{A.134})$$

where S is the number of firms in the S&P 500 index, and where $P_{i,t}$ and $Q_{i,t}$ are the stock price and the number of shares outstanding of firm i at time t , respectively. $LTG_{i,t}$ is the median forecast of firm i 's long term expected earnings growth. The data spans the periods from 1981:12 to 2021:12. All data were downloaded in February 2023.

Finally, I use the difference between survey expectations of log earnings between adjacent years (i.e., $\mathbb{F}_t[e_{t+h-1}]$ and $\mathbb{F}_t[e_{t+h}]$) to obtain $\mathbb{F}_t[\Delta e_{t+h}] = \mathbb{F}_t[e_{t+h}] - \mathbb{F}_t[e_{t+h-1}]$, the time t survey expectation of forward one-year log earnings growth $h = 1, 2, 3, 4$ years ahead. For the $h = 5$ year horizon, I interpret the IBES's Long-Term Growth (LTG) forecast as the 5-year forward annual log earnings growth from 4 to 5 years ahead:

$$\mathbb{F}_t[\Delta e_{t+h}] = \begin{cases} \mathbb{F}_t[e_{t+h}] - \mathbb{F}_t[e_{t+h-1}] & \text{if } h = 1, 2, 3, 4 \text{ years} \\ LTG_t & \text{if } h = 5 \text{ years} \end{cases} \quad (\text{A.135})$$

To estimate any biases in IBES analyst forecasts, the dynamic machine algorithm takes as an input a likely date corresponding to information analysts could have known at the time of their forecast. IBES does not provide an explicit deadline for their forecasts to be returned. Therefore I instead use the “statistical period” day (the day when the set of summary statistics was calculated) as a proxy for the deadline. I set the machine deadline to be the day before this date. The statistical period date is typically between day 14 and day 20 of a given month, implying that the machine deadline varies from month to month. As the machine learning algorithm uses mixed-frequency techniques adapted to quarterly sampling intervals, while the IBES forecasts are monthly, I compare machine and IBES analyst forecasts as of the middle month of each quarter, considering 12-month ahead forecast from the beginning of the month following the survey month.

D.7 Price-Earnings Ratio

I construct a quarterly series for the price-earnings ratio $PE_t \equiv P_t/E_t$ as the ratio between the end-of-quarter S&P 500 stock price index P_t and the S&P 500 quarterly total earnings E_t . I infer subjective expectations of the log price-earnings ratio $\mathbb{F}_t[pe_{t+h}]$ by combining the current log price-earnings ratio pe_t with h year ahead subjective expectations of annual log stock returns $\mathbb{F}_t[r_{t+h}]$ and annual log earnings growth $\mathbb{F}_t[\Delta e_{t+h}]$, following the approach used in De La O and Myers (2021). Rearrange the Campbell and Shiller (1988) log-linear approximation of the price-earnings ratio in equation (A.33) to express the future log price-earnings ratio as a function of current log price-earnings, log earnings growth, and log stock returns:

$$pe_{t+h} = \frac{1}{\rho^h} pe_t - \frac{1}{\rho^h} \sum_{j=1}^h \rho^{j-1} (c_{pe} + \Delta e_{t+j} - r_{t+j}) \quad (\text{A.136})$$

where the equation holds both ex-ante and ex-post. Apply subjective expectations \mathbb{F}_t on both sides of the equation:

$$\mathbb{F}_t[pe_{t+h}] = \frac{1}{\rho^h} pe_t - \frac{1}{\rho^h} \sum_{j=1}^h \rho^{j-1} (c_{pe} + \underbrace{\mathbb{F}_t[\Delta e_{t+j}]}_{\text{Survey (IBES)}} - \underbrace{\mathbb{F}_t[r_{t+j}]}_{\text{Survey (CFO)}}) \quad (\text{A.137})$$

where subjective expectations about j years ahead forward annual log stock returns $\mathbb{F}_t[r_{t+j}]$ and forward annual log earnings growth $\mathbb{F}_t[\Delta e_{t+j}]$ use survey forecasts from the CFO survey and IBES, respectively.

D.8 Earnings-Employment Ratio

The current earnings-employment ratio is defined as $EL_t \equiv E_t/L_{t+1}$, where E_t denotes quarterly total earnings for the S&P 500 and L_{t+1} is the employment stock at the beginning of period $t+1$. I measure L_{t+1} using end-of-period employment levels within each quarter. This timing assumption ensures that the measures are consistent with the timing conventions from Section B while still remaining known to firms by the end of period t . Let $el_t = \log(EL_t)$ denote the log earnings-employment ratio. Subjective expectations of future values $\mathbb{F}_t[el_{t+1}] \equiv \mathbb{F}_t[\log E_{t+1} - \log L_{t+1}]$ are constructed as:

$$\mathbb{F}_t[el_{t+1}] \equiv \mathbb{F}_t[\log E_{t+1} - \log L_{t+1}] \quad (\text{A.138})$$

$$= \mathbb{F}_t[\log E_{t+1} - \log E_t + \log E_t - \log L_{t+1}] \quad (\text{A.139})$$

$$= el_t + \underbrace{\mathbb{F}_t[\Delta e_{t,t+1}]}_{\text{Survey (IBES)}} \quad (\text{A.140})$$

where survey expectations of earnings growth $\mathbb{F}_t[\Delta e_{t,t+h}]$ come from IBES.

D.9 Machine Learning Forecasts

For each survey forecast, I also construct their corresponding machine learning forecast by estimating a Long Short-Term Memory (LSTM) neural network:

$$\mathbb{E}_t[y_{t,t+h}] = G(\mathcal{X}_t, \boldsymbol{\beta}_{h,t}) \quad (\text{A.141})$$

where $y_{t,t+h}$ denotes the variable y to be predicted h years ahead of time t , and \mathcal{X}_t is a large input dataset of right-hand-side variables including the intercept. $G(\mathcal{X}_t, \boldsymbol{\beta}_{h,t})$ denotes predicted values from a LSTM neural network that can be represented by a (potentially) high dimensional set of finite-valued parameters $\boldsymbol{\beta}_{h,t}$. The machine learning model is estimated using an algorithm that takes into account the data-rich environment in which managers operate in (Bianchi et al., 2022 and Bianchi et al., 2024b). When constructing machine learning forecasts of each variable, I allow the machine to use only information that would have been available to all survey respondents at the time of each forecast. See Section E for details about the machine learning algorithm and

predictor variables. Machine expectations about the price-earnings ratio $\mathbb{E}_t[pe_{t+1}]$ and earnings-employment ratio $\mathbb{E}_t[el_{t+1}]$ are constructed similarly to the survey counterpart, by replacing the survey forecasts of stock returns, earnings growth, and employment growth on the right-hand side of equation (A.137) and (19) with the corresponding machine learning forecasts.

E Machine Learning

E.1 Data Inputs for Machine Learning Algorithm

E.1.1 Macro Data Surprises

These data are used as inputs into the machine learning forecasts. I obtain median forecasts for GDP growth (Q/Q percentage change), core CPI (Month/Month change), unemployment rate (percentage point), and nonfarm payroll (month/month change) from the Money Market Service Survey. The median market survey forecasts are compiled and published by the Money Market Services (MMS) the Friday before each release. I apply the approach used in Bauer and Swanson (2023) and define macroeconomic data surprise as the actual value of the data release minus the median expectation from MMS on the Friday immediately prior to that data release. The GDP growth forecasts are available quarterly from 1990Q1 to 2022Q1. The core CPI forecast is available monthly from July 1989 to April 2022. The median forecasts for the unemployment rate and nonfarm payrolls are available monthly from Jan 1980 to May 2022, and Jan. 1985 to May 2022, respectively. All survey forecasts were downloaded from Haver Analytics on December 17, 2022. To pin down the timing of when the news was actually released I follow the published tables of releases from the Bureau of Labor Statistics (BLS), discussed below.

The macro news events are indexed by their date and time of the data release, while the machine learning algorithm is adapted to quarterly sampling frequencies. When including the macro data surprises as additional predictors for the machine forecast, I time-aggregate the macro data surprises to a quarterly frequency by taking the sum of the surprises across data releases that occurred before the response deadline set for the machine. For example, if the response deadline is set to the first day of the middle month of each quarter (e.g., February 1st), I take the sum of the surprises from data releases up to the day before the deadline, the last day of the first month of each quarter (e.g., January 31st).

E.1.2 FOMC Surprises

FOMC surprises are defined as the changes in the current-month, 1, 2, 6, 12, and 24 month-ahead federal funds futures (FFF) contract rate and changes in the 1, 2, 4, and 8 quarter-ahead Eurodollar (ED) futures contract rate, from 10 minutes before to 20 minutes after each U.S. Federal Reserve Federal Open Market Committee (FOMC) announcement. The data on FFF and ED were downloaded on June 3rd 2022. When benchmarking against a survey, I use the last FOMC meeting before the survey deadline to compute surprises. For surveys that do not have a clear deadline, I compute surprises using from the last FOMC in the first month of the

quarter. When benchmarking against moving average, I use the last FOMC meeting before the end of the first month in each quarter to compute surprises.

When including the FOMC surprises as additional predictors for the machine forecast, I time-aggregate the FOMC surprises to a quarterly frequency by taking the sum of the surprises across FOMC announcements that occurred before the response deadline set for the machine. For example, if the response deadline is set to the first day of the middle month of each quarter (e.g., February 1st), I take the sum of the surprises from FOMC announcements up to the day before the deadline, the last day of the first month of each quarter (e.g., January 31st).

E.1.3 S&P 500 Jumps

As a measure of the market’s reaction to news shocks, I use the jump in the S&P 500 pre- and post- a 30-minute window around major news events. The events in our analysis include (i) 1,482 macroeconomic data releases for U.S. GDP, Consumer Price Index (CPI), unemployment, and payroll data spanning 1980:01-2021:12, (ii) 16 corporate earnings announcement days spanning 1999:03-2020:05, and (iii) 219 Federal Open Market Committee (FOMC) press releases from the Fed spanning 1994:02-2021:12. The corporate earnings news events are from Baker et al. (2019) who conduct textual analyses of *Wall Street Journal* articles to identify days in which there were large jumps in the aggregate stock market that could be attributed to corporate earnings news with high confidence.

The jump in the S&P 500 for a given event is defined as

$$j_\tau = p_{\tau+\delta_{post}} - p_{\tau-\delta_{pre}} \quad (\text{A.142})$$

where τ indexes the time of an event and $p_\tau = \log(P_\tau)$ is the log S&P 500 index. δ_{pre} and δ_{post} denote the pre and post event windows, which is 10 minutes before and 20 minutes after the event, respectively.

I obtain data on P_τ using tick-by-tick data on the S&P 500 index from tickdata.com. The series was purchased and downloaded on 7/2/2022 from <https://www.tickdata.com/>. I create the minutely data using the close price within each minute. I supplement the S&P 500 index using S&P500 E-mini futures for events that occur in off-market hours. I use the current-quarter contract futures. I purchased the S&P 500 E-mini futures from CME group on 7/2/2022 at <https://datamine.cmegroup.com/>. Our sample spans 1/2/1986 to 6/30/2022.

For each event, I separate out the events for which the S&P 500 increased over the window ($j_\tau^{(+)} \geq 0$) and those for which the market decreased ($j_\tau^{(-)} \leq 0$):

$$j_\tau^{(+)} = \max\{0, j_\tau\} \quad (\text{A.143})$$

$$j_\tau^{(-)} = \min\{0, j_\tau\} \quad (\text{A.144})$$

I aggregate the event-level jumps to monthly time series by summing over all the relevant events within the month, where the events are partitioned into two groups based on the sign of the

jump:

$$J_t^{(+)} = \sum_{\tau \in x(t)} j_{\tau}^{(+)} \quad (\text{A.145})$$

$$J_t^{(-)} = \sum_{\tau \in x(t)} j_{\tau}^{(-)} \quad (\text{A.146})$$

where t indexes the month and $x(t)$ is the set of all events that occurred within month t . The procedure results in two monthly variables, $J_t^{(+)}$ and $J_t^{(-)}$, which capture total market reaction to news events in either direction during the quarter. The series spans the period 1994:02 to 2022:03. Separating out the events based on the sign of the jump allows us to capture any differential effects on return predictability based on whether the market perceived the news as good or bad. The partition also allows us to accurately capture the total extent of over- or under-reaction. Otherwise, mixing all the events would only capture the net effect of the jumps and bias the market reaction towards zero.

When used as additional predictors in the for the machine forecast, the jumps need to be converted to quarterly time series because the machine learning algorithm is adapted to a quarterly sampling frequency. The set of events in $x(t)$ is chosen so that the machine only sees the news events that would have been available to the real-time manager. When combining the events within a quarter, I impose the response deadline used to produce the machine forecast. For example, if the response deadline is set to the first day of the middle month of each quarter (e.g., February 1st), I use the jumps from the events up to the day before the deadline, the last day of the first month of each quarter (e.g., January 31st).

E.1.4 Real-Time Macro Data

This section gives details on the real time macro data inputs used in the machine learning forecasts. A subset of these series are used in the structural estimation. At each forecast date in the sample, I construct a dataset of macro variables that could have been observed on or before the day of the survey deadline. I use the Philadelphia Fed’s Real-Time Data Set to obtain vintages of macro variables.⁴ These vintages capture changes to historical data due to periodic revisions made by government statistical agencies. The vintages for a particular series can be available at the monthly and/or quarterly frequencies, and the series have monthly and/or quarterly observations. In cases where a variable has both frequencies available for its vintages and/or its observations, I choose one format of the variable. For instance, nominal personal consumption expenditures on goods is quarterly data with both monthly and quarterly vintages available; in this case, I use the version with monthly vintages.

Table A.16 gives the complete list of real-time macro variables. Included in the table is the first available vintages for each variable that has multiple vintages. I do not include the last

⁴The real-time data sets are available at <https://www.philadelphiafed.org/research-and-data/real-time-center/real-time-data/data-files>.

vintage because most variables have vintages through the present.⁵ Table A.16 also lists the transformation applied to each variable to make them stationary before generating factors. Let $X_{i,t}$ denote variable i at time t after the transformation, and let $X_{i,t}^A$ be the untransformed series. Let $\Delta = (1 - L)$ with $LX_{i,t} = X_{i,t-1}$. There are seven possible transformations with the following codes:

- 1 Code lv : $X_{i,t} = X_{i,t}^A$
- 2 Code Δlv : $X_{i,t} = X_{i,t}^A - X_{i,t-1}^A$
- 3 Code $\Delta^2 lv$: $X_{i,t} = \Delta^2 X_{i,t}^A$
- 4 Code ln : $X_{i,t} = \ln(X_{i,t}^A)$
- 5 Code Δln : $X_{i,t} = \ln(X_{i,t}^A) - \ln(X_{i,t-1}^A)$
- 6 Code $\Delta^2 ln$: $X_{i,t} = \Delta^2 \ln(X_{i,t}^A)$
- 7 Code $\Delta lv/lv$: $X_{i,t} = (X_{i,t}^A - X_{i,t-1}^A)/X_{i,t-1}^A$

Table A.16: List of Macro Dataset Variables

No.	Short Name	Source	Tran	Description	First Vintage
Group 1: Output and Income					
1	IPMMVMD	Philly Fed	Δln	Ind. production index - Manufacturing	1962M11
2	IPTMVMD	Philly Fed	Δln	Ind. production index - Total	1962M11
3	CUMMVMD	Philly Fed	lv	Capacity utilization - Manufacturing	1979M8
4	CUTMVMD	Philly Fed	lv	Capacity utilization - Total	1983M7
5	NCPROFATMVQD	Philly Fed	Δln	Nom. corp. profits after tax without IVA/CCAdj	1965Q4
6	NCPROFATWMVQD	Philly Fed	Δln	Nom. corp. profits after tax with IVA/CCAdj	1981Q1
7	OPHMVQD	Philly Fed	Δln	Output per hour - Business sector	1998Q4
8	NDPIQVQD	Philly Fed	Δln	Nom. disposable personal income	1965Q4
9	NOUTPUTQVQD	Philly Fed	Δln	Nom. GNP/GDP	1965Q4
10	NPIQVQD	Philly Fed	Δln	Nom. personal income	1965Q4
11	NPSAVQVQD	Philly Fed	Δlv	Nom. personal saving	1965Q4
12	OLIQVQD	Philly Fed	Δln	Other labor income	1965Q4
13	PINTIQVQD	Philly Fed	Δln	Personal interest income	1965Q4
14	PINTPAIDQVQD	Philly Fed	Δln	Interest paid by consumers	1965Q4
15	PROPIQVQD	Philly Fed	Δln	Proprietors' income	1965Q4
16	PTAXQVQD	Philly Fed	Δln	Personal tax and nontax payments	1965Q4
17	RATESAVQVQD	Philly Fed	Δlv	Personal saving rate	1965Q4
18	RENTIQVQD	Philly Fed	Δlv	Rental income of persons	1965Q4
19	ROUTPUTQVQD	Philly Fed	Δln	Real GNP/GDP	1965Q4
20	SSCONTRIBQVQD	Philly Fed	Δln	Personal contributions for social insurance	1965Q4
21	TRANPFQVQD	Philly Fed	Δln	Personal transfer payments to foreigners	1965Q4
22	TRANRQVQD	Philly Fed	Δln	Transfer payments	1965Q4
23	CUUR0000SA0E	BLS	$\Delta^2 ln$	Energy in U.S. city avg., all urban consumers, not seasonally adj	
Group 2: Employment					
24	EMPLOYMVMD	Philly Fed	Δln	Nonfarm payroll	1946M12

⁵For variables BASEBASAQVMD, NBRBASAQVMD, NBRECBASAQVMD, and TRBASAQVMD, the last available vintage is 2013Q2.

No.	Short Name	Source	Tran	Description	First Vintage
25	HMVMD	Philly Fed	lv	Aggregate weekly hours - Total	1971M9
26	HGMVMD	Philly Fed	lv	Agg. weekly hours - Goods-producing	1971M9
27	HSMVMD	Philly Fed	lv	Agg. weekly hours - Service-producing	1971M9
28	LFCMVMD	Philly Fed	Δln	Civilian labor force	1998M11
29	LFPARTMVMD	Philly Fed	lv	Civilian participation rate	1998M11
30	POPMVMD	Philly Fed	Δln	Civilian noninstitutional population	1998M11
31	ULCMVQD	Philly Fed	Δln	Unit labor costs - Business sector	1998Q4
32	RUCQVMD	Philly Fed	Δlv	Unemployment rate	1965Q4
33	WSDQVQD	Philly Fed	Δln	Wage and salary disbursements	1965Q4
Group 3: Orders, Investment, Housing					
34	HSTARTSMVMD	Philly Fed	Δln	Housing starts	1968M2
35	RINVBFMVQD	Philly Fed	Δln	Real gross private domestic inv. - Nonresidential	1965Q4
36	RINVCHIMVQD	Philly Fed	Δlv	Real gross private domestic inv. - Change in private inventories	1965Q4
37	RINVRESIDMVQD	Philly Fed	Δln	Real gross private domestic inv. - Residential	1965Q4
38	CASESHILLER	S&P	Δln	Case-Shiller US National Home Price index/CPI	1987M1
Group 4: Consumption					
39	NCONGMMVMD	Philly Fed	Δln	Nom. personal cons. exp. - Goods	2009M8
40	NCONHHMMVMD	Philly Fed	Δln	Nom. hh. cons. exp.	2009M8
41	NCONSHHMMVMD	Philly Fed	Δln	Nom. hh. cons. exp. - Services	2009M8
42	NCONSNPMMVMD	Philly Fed	Δln	Nom. final cons. exp. of NPISH	2009M8
43	RCONDMMVMD	Philly Fed	Δln	Real personal cons. exp. - Durables	1998M11
44	RCONGMMVMD	Philly Fed	Δln	Real personal cons. exp. - Goods	2009M8
45	RCONHHMMVMD	Philly Fed	Δln	Real hh. cons. exp.	2009M8
46	RCONMMVMD	Philly Fed	Δln	Real personal cons. exp. - Total	1998M11
47	RCONNDMVMD	Philly Fed	Δln	Real personal cons. exp. - Nondurables	1998M11
48	RCONSHHMMVMD	Philly Fed	Δln	Real hh. cons. exp. - Services	2009M8
49	RCONSMVMD	Philly Fed	Δln	Real personal cons. exp. - Services	1998M11
50	RCONSNPMMVMD	Philly Fed	Δln	Real final cons. exp. of NPISH	2009M8
51	NCONGMVQD	Philly Fed	Δln	Nom. personal cons. exp. - Goods	2009Q3
52	NCONHHMVQD	Philly Fed	Δln	Nom. hh. cons. exp.	2009Q3
53	NCONSHHMVQD	Philly Fed	Δln	Nom. hh. cons. exp. - Services	2009Q3
54	NCONSNPMVQD	Philly Fed	Δln	Nom. final cons. exp. of NPISH	2009Q3
55	RCONDMVQD	Philly Fed	Δln	Real personal cons. exp. - Durable goods	1965Q4
56	RCONGMVQD	Philly Fed	Δln	Real personal cons. exp. - Goods	2009Q3
57	RCONHHMVQD	Philly Fed	Δln	Real hh. cons. exp.	2009Q3
58	RCONMVQD	Philly Fed	Δln	Real personal cons. exp. - Total	1965Q4
59	RCONNDMVQD	Philly Fed	Δln	Real personal cons. exp. - Nondurable goods	1965Q4
60	RCONSHHMVQD	Philly Fed	Δln	Real hh. cons. exp. - Services	2009Q3
61	RCONSMVQD	Philly Fed	Δln	Real personal cons. exp. - Services	1965Q4
62	RCONSNPMVQD	Philly Fed	Δln	Real final cons. exp. of NPISH	2009Q3
63	NCONQVQD	Philly Fed	Δln	Nom. personal cons. exp.	1965Q4
Group 5: Prices					
64	PCONGMMVMD	Philly Fed	$\Delta^2 ln$	Price index for personal cons. exp. - Goods	2009M8
65	PCONHHMMVMD	Philly Fed	$\Delta^2 ln$	Price index for hh. cons. exp.	2009M8
66	PCONSHHMMVMD	Philly Fed	$\Delta^2 ln$	Price index for hh. cons. exp. - Services	2009M8
67	PCONSNPMMVMD	Philly Fed	$\Delta^2 ln$	Price index for final cons. exp. of NPISH	2009M8
68	PCPIMVMD	Philly Fed	$\Delta^2 ln$	Consumer price index	1998M11
69	PCPIXMVMD	Philly Fed	$\Delta^2 ln$	Core consumer price index	1998M11
70	PPPIMVMD	Philly Fed	$\Delta^2 ln$	Producer price index	1998M11
71	PPPIXMVMD	Philly Fed	$\Delta^2 ln$	Core producer price index	1998M11
72	PCONGMVQD	Philly Fed	$\Delta^2 ln$	Price index for personal. cons. exp. - Goods	2009Q3
73	PCONHHMVQD	Philly Fed	$\Delta^2 ln$	Price index for hh. cons. exp.	2009Q3

No.	Short Name	Source	Tran	Description	First Vintage
74	PCONSHHMOVQD	Philly Fed	$\Delta^2 \ln$	Price index for hh. cons. exp. - Services	2009Q3
75	PCONSNNMOVQD	Philly Fed	$\Delta^2 \ln$	Price index for final cons. exp. of NPISH	2009Q3
76	PCONXMOVQD	Philly Fed	$\Delta^2 \ln$	Core price index for personal cons. exp.	1996Q1
77	CPIQVMD	Philly Fed	$\Delta^2 \ln$	Consumer price index	1994Q3
78	PQVQD	Philly Fed	$\Delta^2 \ln$	Price index for GNP/GDP	1965Q4
79	PCONQVQD	Philly Fed	$\Delta^2 \ln$	Price index for personal cons. exp.	1965Q4
80	PIMPQVQD	Philly Fed	$\Delta^2 \ln$	Price index for imports of goods and services	1965Q4
Group 6: Trade and Government					
81	REXMOVQD	Philly Fed	$\Delta \ln$	Real exports of goods and services	1965Q4
82	RGMVQD	Philly Fed	$\Delta \ln$	Real government cons. and gross inv. - Total	1965Q4
83	RGFMVQD	Philly Fed	$\Delta \ln$	Real government cons. and gross inv. - Federal	1965Q4
84	RGSLMVQD	Philly Fed	$\Delta \ln$	Real government cons. and gross. inv. - State and local	1965Q4
85	RIMPMVQD	Philly Fed	$\Delta \ln$	Real imports of goods and services	1965Q4
86	RNXMOVQD	Philly Fed	$\Delta \ln$	Real net exports of goods and services	1965Q4
Group 7: Money and Credit					
87	BASEBASAQVMD	Philly Fed	$\Delta^2 \ln$	Monetary base	1980Q2
88	M1QVMD	Philly Fed	$\Delta^2 \ln$	M1 money stock	1965Q4
89	M2QVMD	Philly Fed	$\Delta^2 \ln$	M2 money stock	1971Q2
90	NBRBASAQVMD	Philly Fed	$\Delta \ln / \ln$	Nonborrowed reserves	1967Q3
91	NBRECASAQVMD	Philly Fed	$\Delta \ln / \ln$	Nonborrowed reserves plus extended credit	1984Q2
92	TRBASAQVMD	Philly Fed	$\Delta^2 \ln$	Total reserves	1967Q3
93	DIVQVQD	Philly Fed	$\Delta \ln$	Dividends	1965Q4

E.1.5 Daily Financial Data

Daily Data and construction of daily factors These data are used in the machine learning forecasts. The daily financial series in this data set are from the daily financial dataset used in Andreou et al. (2013). I create a smaller daily database which is a subset of the large cross-section of 991 daily series in their dataset. Our dataset covers five classes of financial assets: (i) the Commodities class; (ii) the Corporate Risk category; (iii) the Equities class; (iv) the Foreign Exchange Rates class and (v) the Government Securities.

The dataset includes up to 87 daily predictors in a daily frequency from 23-Oct-1959 to 24-Oct-2021 (14852 trading days) from the above five categories of financial assets. I remove series with fewer than ten years of data and time periods with no variables observed, which occurs for some series in the early part of the sample. For those years, I have less than 87 series. There are 39 commodity variables which include commodity indices, prices and futures, 16 corporate risk series, 9 equity series which include major US stock market indices and the 500 Implied Volatility, 16 government securities which include the federal funds rate, government treasury bills of securities from three months to ten years, and 7 foreign exchange variables which include the individual foreign exchange rates of major five US trading partners and two effective exchange rate. I choose these daily predictors because they are proposed in the literature as good predictors of economic growth.

I construct daily financial factors in a quarterly frequency in two steps. First, I use these daily financial time series to form factors at a daily frequency. The raw data used to form factors are

always transformed to achieve stationarity and standardized before performing factor estimation (see generic description below). I re-estimate factors at each date in the sample recursively over time using the entire history of data available in real time prior to each out-of-sample forecast.

In the second step, I convert these daily financial indicators to quarterly weighted variables to form quarterly factors by selecting an optimal weighting scheme according to the method described below (see the weighting scheme section).

The data series used in this dataset are listed below in Table A.17 by data source. The tables also list the transformation applied to each variable to make them stationary before generating factors. The transformations used to stationarize a time series are the same as those explained in the section “Monthly financial factor data”.

Table A.17: List of Daily Financial Dataset Variables

No.	Short Name	Source	Tran	Description
Group 1: Commodities				
1	GSIZSPT	Data Stream	$\Delta \ln$	S&P GSCI Zinc Spot - PRICE INDEX
2	GSSBSPT	Data Stream	$\Delta \ln$	S&P GSCI Sugar Spot - PRICE INDEX
3	GSSOSPT	Data Stream	$\Delta \ln$	S&P GSCI Soybeans Spot - PRICE INDEX
4	GSSISPT	Data Stream	$\Delta \ln$	S&P GSCI Silver Spot - PRICE INDEX
5	GSIKSPT	Data Stream	$\Delta \ln$	S&P GSCI Nickel Spot - PRICE INDEX
6	GSLCSPT	Data Stream	$\Delta \ln$	S&P GSCI Live Cattle Spot - PRICE INDEX
7	GSLHSPT	Data Stream	$\Delta \ln$	S&P GSCI Lean Hogs Index Spot - PRICE INDEX
8	GSILSPT	Data Stream	$\Delta \ln$	S&P GSCI Lead Spot - PRICE INDEX
9	GSGCSPT	Data Stream	$\Delta \ln$	S&P GSCI Gold Spot - PRICE INDEX
10	GSCTSPT	Data Stream	$\Delta \ln$	S&P GSCI Cotton Spot - PRICE INDEX
11	GSKCSPT	Data Stream	$\Delta \ln$	S&P GSCI Coffee Spot - PRICE INDEX
12	GSCCSPT	Data Stream	$\Delta \ln$	S&P GSCI Cocoa Index Spot - PRICE INDEX
13	GSIASPT	Data Stream	$\Delta \ln$	S&P GSCI Aluminum Spot - PRICE INDEX
14	SGWTSPT	Data Stream	$\Delta \ln$	S&P GSCI All Wheat Spot - PRICE INDEX
15	EIAEBRT	Data Stream	$\Delta \ln$	Europe Brent Spot FOB U\$/BBL Daily
16	CRUDOIL	Data Stream	$\Delta \ln$	Crude Oil-WTI Spot Cushing U\$/BBL - MID PRICE
17	LTICASH	Data Stream	$\Delta \ln$	LME-Tin 99.85% Cash U\$/MT
18	CWFCS00	Data Stream	$\Delta \ln$	CBT-WHEAT COMPOSITE FUTURES CONT. - SETT. PRICE
19	CCFCS00	Data Stream	$\Delta \ln$	CBT-CORN COMP. CONTINUOUS - SETT. PRICE
20	CSYCS00	Data Stream	$\Delta \ln$	CBT-SOYBEANS COMP. CONT. - SETT. PRICE
21	NCTCS20	Data Stream	$\Delta \ln$	CSCE-COTTON #2 CONT.2ND FUT - SETT. PRICE
22	NSBCS00	Data Stream	$\Delta \ln$	CSCE-SUGAR #11 CONTINUOUS - SETT. PRICE
23	NKCCS00	Data Stream	$\Delta \ln$	CSCE-COFFEE C CONTINUOUS - SETT. PRICE
24	NCCCS00	Data Stream	$\Delta \ln$	CSCE-COCOA CONTINUOUS - SETT. PRICE
25	CZLCS00	Data Stream	$\Delta \ln$	ECBOT-SOYBEAN OIL CONTINUOUS - SETT. PRICE
26	COFC01	Data Stream	$\Delta \ln$	CBT-OATS COMP. TRc1 - SETT. PRICE
27	CLDCS00	Data Stream	$\Delta \ln$	CME-LIVE CATTLE COMP. CONTINUOUS - SETT. PRICE
28	CLGC01	Data Stream	$\Delta \ln$	CME-LEAN HOGS COMP. TRc1 - SETT. PRICE
29	NGCCS00	Data Stream	$\Delta \ln$	CMX-GOLD 100 OZ CONTINUOUS - SETT. PRICE
30	LAH3MTH	Data Stream	$\Delta \ln$	LME-Aluminium 99.7% 3 Months U\$/MT
31	LED3MTH	Data Stream	$\Delta \ln$	LME-Lead 3 Months U\$/MT
32	LNI3MTH	Data Stream	$\Delta \ln$	LME-Nickel 3 Months U\$/MT
33	LTi3MTH	Data Stream	$\Delta \ln$	LME-Tin 99.85% 3 Months U\$/MT

No.	Short Name	Source	Tran	Description
34	PLNYD	www.macrotrends.net	Δln	Platinum Cash Price (US\$ per troy ounce)
35	XPDD	www.macrotrends.net	Δln	Palladium (US\$ per troy ounce)
36	CUS2D	www.macrotrends.net	Δln	Corn Spot Price (US\$/Bushel)
37	SoybOil	www.macrotrends.net	Δln	Soybean Oil Price (US\$/Pound)
38	OATSD	www.macrotrends.net	Δln	Oat Spot Price (US\$/Bushel)
39	WTIOilFut	US EIA	Δln	Light Sweet Crude Oil Futures Price: 1St Expiring Contract Settlement (\$/Bbl)
Group 2: Equities				
40	S&PCOMP	Data Stream	Δln	S&P 500 COMPOSITE - PRICE INDEX
41	ISPCS00	Data Stream	Δln	CME-S&P 500 INDEX CONTINUOUS - SETT. PRICE
42	SP5EIND	Data Stream	Δln	S&P500 ES INDUSTRIALS - PRICE INDEX
43	DJINDUS	Data Stream	Δln	DOW JONES INDUSTRIALS - PRICE INDEX
44	CYMCS00	Data Stream	Δln	CBT-MINI DOW JONES CONTINUOUS - SETT. PRICE
45	NASCOMP	Data Stream	Δln	NASDAQ COMPOSITE - PRICE INDEX
46	NASA100	Data Stream	Δln	NASDAQ 100 - PRICE INDEX
47	CBOEVIX	Data Stream	lv	CBOE SPX VOLATILITY VIX (NEW) - PRICE INDEX
48	S&P500toVIX	Data Stream	Δln	S&P500/VIX
Group 3: Corporate Risk				
49	LIBOR	FRED	Δlv	Overnight London Interbank Offered Rate (%)
50	1MLIBOR	FRED	Δlv	1-Month London Interbank Offered Rate (%)
51	3MLIBOR	FRED	Δlv	3-Month London Interbank Offered Rate (%)
52	6MLIBOR	FRED	Δlv	6-Month London Interbank Offered Rate (%)
53	1YLIBOR	FRED	Δlv	One-Year London Interbank Offered Rate (%)
54	1MEuro-FF	FRED	lv	1-Month Eurodollar Deposits (London Bid) (% P.A.) minus Fed Funds
55	3MEuro-FF	FRED	lv	3-Month Eurodollar Deposits (London Bid) (% P.A.) minus Fed Funds
56	6MEuro-FF	FRED	lv	6-Month Eurodollar Deposits (London Bid) (% P.A.) minus Fed Funds
57	APFNF-AANF	Data Stream	lv	1-Month A2/P2/F2 Nonfinancial Commercial Paper (NCP) (% P. A.) minus 1-Month Aa NCP (% P.A.)
58	APFNF-AAF	Data Stream	lv	1-Month A2/P2/F2 NCP (% P.A.) minus 1-Month Aa Financial Commercial Paper (% P.A.)
59	TED	Data Stream, FRED	lv	3Month Tbill minus 3-Month London Interbank Offered Rate (%)
60	MAaa-10YTB	Data Stream	lv	Moody Seasoned Aaa Corporate Bond Yield (% P.A.) minus Y10-Tbond
61	MBaa-10YTB	Data Stream	lv	Moody Seasoned Baa Corporate Bond Yield (% P.A.) minus Y10-Tbond
62	MLA-10YTB	Data Stream, FRED	lv	Merrill Lynch Corporate Bonds: A Rated: Effective Yield (%) minus Y10-Tbond
63	MLAA-10YTB	Data Stream, FRED	lv	Merrill Lynch Corporate Bonds: Aa Rated: Effective Yield (%) minus Y10-Tbond
64	MLAAA-10YTB	Data Stream, FRED	lv	Merrill Lynch Corporate Bonds: Aaa Rated: Effective Yield (%) minus Y10-Tbond
Group 4: Treasuries				
65	FRFEDFD	Data Stream	Δlv	US FED FUNDS EFF RATE (D) - MIDDLE RATE
66	FRTBS3M	Data Stream	Δlv	US T-BILL SEC MARKET 3 MONTH (D) - MIDDLE RATE
67	FRTBS6M	Data Stream	Δlv	US T-BILL SEC MARKET 6 MONTH (D) - MIDDLE RATE
68	FRTCM1Y	Data Stream	Δlv	US TREASURY CONST MAT 1 YEAR (D) - MIDDLE RATE
69	FRTCM10	Data Stream	Δlv	US TREASURY CONST MAT 10 YEAR (D) - MIDDLE RATE

No.	Short Name	Source	Tran	Description
70	6MTB-FF	Data Stream	lv	6-month treasury bill market bid yield at constant maturity (%) minus Fed Funds
71	1YTB-FF	Data Stream	lv	1-year treasury bill yield at constant maturity (% P.A.) minus Fed Funds
72	10YTB-FF	Data Stream	lv	10-year treasury bond yield at constant maturity (% P.A.) minus Fed Funds
73	6MTB-3MTB	Data Stream	lv	6-month treasury bill yield at constant maturity (% P.A.) minus 3M-Tbills
74	1YTB-3MTB	Data Stream	lv	1-year treasury bill yield at constant maturity (% P.A.) minus 3M-Tbills
75	10YTB-3MTB	Data Stream	lv	10-year treasury bond yield at constant maturity (% P.A.) minus 3M-Tbills
76	BKEVEN05	FRB	lv	US Inflation compensation: continuously compounded zero-coupon yield: 5-year (%)
77	BKEVEN10	FRB	lv	US Inflation compensation: continuously compounded zero-coupon yield: 10-year (%)
78	BKEVEN1F4	FRB	lv	BKEVEN1F4
79	BKEVEN1F9	FRB	lv	BKEVEN1F9
80	BKEVEN5F5	FRB	lv	US Inflation compensation: coupon equivalent forward rate: 5-10 years (%)
Group 5: Foreign Exchange (FX)				
81	US_CWBN	Data Stream	Δln	US NOMINAL DOLLAR BROAD INDEX - EXCHANGE INDEX
82	US_CWMN	Data Stream	Δln	US NOMINAL DOLLAR MAJOR CURR INDEX - EXCHANGE INDEX
83	US_CSFR2	Data Stream	Δln	CANADIAN \$ TO US \$ NOON NY - EXCHANGE RATE
84	EU_USFR2	Data Stream	Δln	EURO TO US\$ NOON NY - EXCHANGE RATE
85	US_YFR2	Data Stream	Δln	JAPANESE YEN TO US \$ NOON NY - EXCHANGE RATE
86	US_SFFR2	Data Stream	Δln	SWISS FRANC TO US \$ NOON NY - EXCHANGE RATE
87	US_UKFR2	Data Stream	Δln	UK POUND TO US \$ NOON NY - EXCHANGE RATE

E.1.6 LDA Data

The LDA data are used as inputs into the machine learning forecasts. The database for our Latent Dirichlet Allocation (LDA) analysis contains around one million articles published in *Wall Street Journal* between January 1984 to June 2022. The current vintage of the results reported here is based a randomly selected sub-sample of 200,000 articles over the same period, one-fifth size of the entire database. The sample selection procedures follows Bybee et al. (2021). First, I remove all articles prior to January 1984 and after June 2022 and exclude articles published in weekends. Second, I exclude articles with subject tags associated with obviously non-economic content such as sports. Third, I exclude articles with the certain headline patterns, such as those associated with data tables or those corresponding to regular sports, leisure, or books columns. I filter the articles using the same list of exclusions provided by Bybee et al. (2021). Last, I exclude articles with less than 100 words.

Processing of texts The processing of the texts can be summarized into five steps:

1. Tokenization: parse each article’s text into a white-space-separated word list retaining the article’s word ordering.
2. I drop all non-alphabetical characters and set the remaining characters to lower-case, remove words with less than 3 letters, and remove common stop words and URL-based terms. I use a standard list of stop words from the Python library *gensim.parsing.preprocessing*.
3. Lemmatization and Stemming: lemmatization returns the original form of a word using external dictionary *Textblob.Word* in Python and based on the context of the word. For instance, as a verb, “went” is converted to “go”. Stemming usually refers to a heuristic process that remove the trailing letters at the end of the words, such as from “assesses” to “assess”, and “really” to “real”. I use the Python library *Textblob.Word* to implement the lemmatization and *SnowballStemmer* for the stemming. The results are not very sensitive to the particular Python packages being used.
4. From the first three steps, I obtain a list of uni-grams which are a list of singular words. For example, “united” and “states” are uni-grams from “united states”. From the list of uni-grams, I generate a set of bi-grams as all pairs of (ordered) adjacent uni-grams. For example, “united states” together is one bi-gram. I then exclude uni-grams and bi-grams appearing in less than 0.1% of articles.
5. Last, I convert an article’s word list into a vector of counts for each uni-gram and bi-gram. For example, the vector of counts [5, 7, 2] corresponds to the number of times the words [“federal”, “reserve”, “bank”] appear in the article.

The LDA Model The LDA model Blei et al. (2003) essentially achieves substantial dimension reduction of the word distribution of each article using the following assumptions. I assume a factor structure on the vectors of word counts. Each factor is a topic and each article is a parametric distribution of topics, specified as follows,

$$\underbrace{\overbrace{w_i}^{V \times 1}}_{\text{word dist of article } i} \sim \text{Mult} \left(\underbrace{\overbrace{\Phi'}^{V \times K}}_{\text{topic-word dist.}}, \underbrace{\overbrace{\theta_i}^{K \times 1}}_{\text{topic dist.}}, \underbrace{N_i}_{\# \text{ of words}} \right) \quad (\text{A.147})$$

where Mult is the multinomial distribution. In the above equation, w_i is a vector of word counts of each unique term (uni-gram or bi-gram) in article i , whose size is equal to the number of unique terms V . K is the number of factors in article i . In the estimation, I assume $K = 180$ following Bybee et al. (2021). Φ is a matrix sized $K \times V$, whose k th row and v th column is equal to the probability of the unique term v showing up in topic k . θ_i stores the weights of all k topics contained in article i , which sum up to one. Dimension reduction is achieved as long as $K \ll V$ (the number of topics are significantly smaller than the number of unique terms).

More specifically, it reduces the dimension from $T \times V$ to $T \times K$ (the size of θ) + $K \times V$ (the size of Φ).

Real-time news factors. I also generate real-time news factors for each month t starting from January 1991. In theory, I could train the LDA model using each real-time monthly vintage but it is computationally challenging. Instead, I simplify the procedure by training the LDA model using quarterly vintages $t, t+3, t+6$, etc, and use the LDA model parameters estimated at t to filter news paper articles within the quarter and generate news factors for those months. More specifically, given every article's word distribution $w_{i,t+s}$, for $s = 0, 1, 2$, and the estimated real-time topic-word distribution parameters $\hat{\Phi}_t$ using articles till date t , one can obtain the filtered topic distribution of each article $\hat{\theta}_{i,t+s}$, as follows,

$$\underbrace{w_{i,t+s}}_{\substack{V \times 1 \\ \text{word dist of article } i \text{ at time } t+s}} \sim \text{Mult} \left(\underbrace{\hat{\Phi}'}_{\substack{V \times K \\ \text{topic-word dist.}}}, \underbrace{\hat{\theta}_{i,t+s}}_{\substack{K \times 1 \\ \text{topic dist.}}}, \underbrace{N_{i,t+s}}_{\substack{\# \text{ of words}}} \right). \quad (\text{A.148})$$

LDA Estimation I use the built-in LDA model estimation toolbox in the Python library <https://pypi.org/project/gensim/Gensim> to implement the model estimation. The model requires following initial inputs and parameters and it is estimated using Bayesian methods.⁶

1. I create a document-term matrix \mathbf{W} as a collection of w_i for all articles i in the sample. The number of rows in \mathbf{W} is equal to the number of articles in our sample and the number of columns in \mathbf{W} is equal to the number of unique uni-gram and bi-grams (after being filtered) across all articles. The matrix \mathbf{W} is used as an input for the LDA model estimation. I then follow Bybee et al. (2021) and set the number of topics K to be 180.⁷
2. In the Python library Gensim, the key parameters of the LDA estim are α and β . With a higher value of α , the documents are composed of more topics. With a higher values of β , each topic contains more terms (uni- or bi-grams). In the implementations, I do not impose any explicit restrictions on initial values of those parameters and set them to be "auto". These two parameters, alongside Φ' and $\{\theta_i\}_i$, are estimated by the toolbox from Python library <https://pypi.org/project/gensim/Gensim>.

Real-time LDA Factors With the estimated topic weights $\theta_{i,t}$ of each article i from the LDA model, I fruther construct time series of the overall news attention to each topic, or a news factor. The value of the topic k at time t is the average weights of topic k of all articles published at t , specified as follows,

$$F_{k,t} = \frac{\sum_i \hat{\theta}_{i,k,t}}{\# \text{ of articles at } t} \quad (\text{A.149})$$

for all topics k .

⁶In theory, maximum-likelihood estimation is possible but it is computationally challenging.

⁷The authors used Bayesian criteria to find 180 to be an optimal number of topics.

E.2 Machine Algorithm Details

The basic dynamic algorithm follows the six step approach of Bianchi et al. (2022) of 1. Sample partitioning, 2. In-sample estimation, 3. Training and cross-validation, 4. Grid reoptimization, 5. Out-of-sample prediction, and 6. Roll forward and repeat. We refer the interested reader to that paper for details and discuss details of the implementation here only insofar as they differ.

At time t , a prior sample of size \dot{T} is partitioned into two subsample windows: a *training sample* consisting of the first T_E observations, and a hold-out *validation sample* of T_V subsequent observations so that $\dot{T} = T_E + T_V$. The training sample is used to estimate the model subject to a specific set of tuning parameter values, and the validation sample is used for tuning the hyperparameters. The model to be estimated over the training sample is

$$y_{t,t+h} = G^e(\mathcal{X}_t, \boldsymbol{\beta}_{h,t}) + \epsilon_{t+h}.$$

where $y_{t,t+h}$ is a time series indexed by j whose value in period $h \geq 1$ the machine is asked to predict at time t , \mathcal{X}_t is a large input dataset of right-hand-side variables including the intercept, and $G^e(\cdot)$ is a machine learning estimator that can be represented by a (potentially) high dimensional set of finite-valued parameters $\boldsymbol{\beta}_{h,t}^e$. We consider two estimators for $G^e(\cdot)$: Elastic Net $G^{\text{EN}}(\mathcal{X}_t, \boldsymbol{\beta}_{j,h}^{\text{EN}})$, and Long Short-Term Memory (LSTM) network $G^{\text{LSTM}}(\mathcal{X}_t, \boldsymbol{\beta}_{j,h}^{\text{LSTM}})$. The $e \in \{\text{EN}, \text{LSTM}\}$ superscripts on $\boldsymbol{\beta}$ indicate that the parameters depend on the estimator being used (See the next section for a description of EN and LSTM). \mathcal{X}_t always denotes the most recent data that would have been in real time prior to the date on which the forecast was submitted. To ensure that the effect of each variable in the input vector is regularized fairly during the estimation, we standardize the elements of \mathcal{X}_t such that sample means are zero and sample standard deviations are unity. It should be noted that the most recent observation on the left-hand-side is generally available in real time only with a one-period lag, thus the forecasting estimations can only be run with data over a sample that stops one period later than today in real time.

The parameters $\boldsymbol{\beta}_{h,t}^e$ are estimated by minimizing the mean-square loss function over the training sample with L_1 and L_2 penalties

$$L(\boldsymbol{\beta}_{h,t}^e, \mathbf{X}_{T_E}, \boldsymbol{\lambda}_t^e) \equiv \underbrace{\frac{1}{T_E} \sum_{\tau=1}^{T_E} (y_{\tau+h} - G^e(\mathcal{X}_\tau, \boldsymbol{\beta}_{h,t}^e))^2}_{\text{Mean Square Error}} + \underbrace{\lambda_{1,t}^e \sum_{k=1}^K |\boldsymbol{\beta}_{j,h,t,k}^e|}_{L_1 \text{ Penalty}} + \underbrace{\lambda_{2,t}^e \sum_{k=1}^K (\boldsymbol{\beta}_{j,h,t,k}^e)^2}_{L_2 \text{ Penalty}}$$

where $\mathbf{X}_{T_E} = (y_{t-T_E}, \dots, y_t, \mathcal{X}'_{t-T_E}, \dots, \mathcal{X}'_t)'$ is the vector containing all observations in the training sample of size T_E . The estimated $\boldsymbol{\beta}_{h,t}^e$ is a function of the data \mathbf{X}_{T_E} and a non-negative regularization parameter vector $\boldsymbol{\lambda}_t^e = (\lambda_{1,t}^e, \lambda_{2,t}^e, \boldsymbol{\lambda}_{0,t}^{\text{LSTM}})'$ where $\boldsymbol{\lambda}_{0,t}^{\text{LSTM}}$ is a set of hyperparameters only relevant when using the LSTM estimator for $G^e(\cdot)$ (see below). For the EN case there are only two hyperparameters, which determine the optimal shrinkage and sparsity of the time t machine specification. The regularization parameters $\boldsymbol{\lambda}_t^e$ are estimated by minimizing the mean-square loss over pseudo-out-of-sample forecast errors generated from rolling regressions through

the validation sample:

$$\begin{aligned}
\hat{\lambda}_t^{EN}, \hat{T}_E, \hat{T}_V = \underset{\lambda_t^{EN}, T_E, T_V}{\operatorname{argmin}} & \left\{ \frac{1}{T_V - h} \sum_{\tau=T_E}^{T_E+T_V-h} \left(y_{\tau+h} - G^{EN}(\mathcal{X}_\tau, \hat{\beta}_{j,h,\tau}^{EN}(\mathbf{X}_{T_E})) \right)^2 \right. \\
& \left. + \underbrace{\lambda_{1,t}^{EN} \sum_{k=1}^K \beta_{j,h,t,k}^{EN}}_{L_1 \text{ Penalty}} + \underbrace{\lambda_{2,t}^{EN} \sum_{k=1}^K (\beta_{j,h,t,k}^{EN})^2}_{L_2 \text{ Penalty}} \right\} \\
\hat{\lambda}_t^{LSTM}, \hat{T}_E, \hat{T}_V = \underset{\lambda_t^{LSTM}, T_E, T_V}{\operatorname{argmin}} & \left\{ \frac{1}{T_V - h} \sum_{\tau=T_E}^{T_E+T_V-h} \left(y_{\tau+h} - G^{LSTM}(\mathcal{X}_\tau, \hat{\beta}_{j,h,\tau}^{LSTM}(\mathbf{X}_{T_E}, \lambda_t^{LSTM})) \right)^2 \right. \\
& \left. + \underbrace{\lambda_{1,t}^{LSTM} \sum_{k=1}^K |\beta_{j,h,t,k}^{LSTM}|}_{L_1 \text{ Penalty}} + \underbrace{\lambda_{2,t}^{LSTM} \sum_{k=1}^K (\beta_{j,h,t,k}^{LSTM})^2}_{L_2 \text{ Penalty}} \right\}
\end{aligned}$$

where $\hat{\beta}_{j,h,\tau}^e(\cdot)$ for $e \in \{EN, LSTM\}$ is the time τ estimate of $\beta_{j,h}^e$ given λ_t^e and data through time τ in a training sample of size T_E . Denote the combined final estimator $\hat{\beta}_{h,t}^e(\mathbf{X}_{\hat{T}_E}, \hat{\lambda}_t^e)$, where the regularization parameter $\hat{\lambda}_t^e$ is estimated using cross-validation dynamically over time. Note that the algorithm also asks the machine to dynamically choose both the optimal training window \hat{T}_E and the optimal validation window \hat{T}_V by minimizing the pseudo-out-of-sample MSE.

The estimation of $\hat{\beta}_{h,t}^e(\mathbf{X}_{\hat{T}_E}, \hat{\lambda}_t^e)$ is repeated sequentially in rolling subsamples, with parameters estimated from information known at time t . Note that the time t subscripts of $\hat{\beta}_{h,t}^e$ and $\hat{\lambda}_t^e$ denote one in a sequence of time-invariant parameter estimates obtained from rolling subsamples, rather than estimates that vary over time within a sample. Likewise, we denote the time t machine belief about $y_{t,t+h}$ as $\mathbb{E}_t^e[y_{t,t+h}]$, defined by

$$\mathbb{E}_t^e[y_{t,t+h}] \equiv G^e \left(\mathcal{X}_t, \hat{\beta}_{h,t}^e(\mathbf{X}_{\hat{T}_E}, \hat{\lambda}_t^e) \right)$$

Finally, the machine MSE is computed by averaging across the sequence of squared forecast errors in the true out-of-sample forecasts for periods $t = (\hat{T} + h), \dots, T$ where T is the last period of our sample. The true out-of-sample forecasts used for neither estimation nor tuning is the *testing subsample* used to evaluate the model's predictive performance.

On rare occasions, one or more of the explanatory variables used in the machine forecast specification assumes a value that is order of magnitudes different from its historical value. This is usually indicative of a measurement problem in the raw data. We therefore program the machine to detect in real-time whether its forecast is an extreme outlier, and in that case to discard the forecast replacing it with the historical mean. Specifically, at each t , the machine forecast $\mathbb{E}_t^e[y_{t,t+h}]$ is set to be the historical mean calculated up to time t whenever the former is five or more standard deviations above its own rolling mean over the most recent 20 quarters.

We include the contemporaneous survey forecasts $\mathbb{F}_t[y_{t,t+h}]$ for the median respondent only for inflation and GDP forecasts, following Bianchi et al. (2022). This procedure allows the machine to capture intangible information due to judgement or private signals. Specifically, for these

forecasts of inflation and GDP growth, we consider the following machine learning empirical specification for forecasting $y_{t,t+h}$ given information at time t , to be benchmarked against the time t survey forecast of respondent-type X , where this type is the median here:

$$y_{t,t+h} = G_{jh}^e(\mathbf{Z}_t) + \gamma_{jh\mathbb{M}} \mathbb{F}_t[y_{t,t+h}] + \epsilon_{t+h}, \quad h \geq 1 \quad (\text{A.150})$$

where $\gamma_{jh\mathbb{M}}$ is a parameter to be estimated, and where $G_{jh\mathbb{M}}(\mathbf{Z}_t)$ represents a ML estimator as function of big data. Note that the intercept α_{jh} from Bianchi et al. (2022) gets absorbed into the $G_{jh}^e(\mathbf{Z}_t)$ in LSTM via the outermost bias term.

E.2.1 Elastic Net (EN)

We use the Elastic Net (EN) estimator, which combines Least Absolute Shrinkage and Selection Operator (LASSO) and ridge type penalties. The model can be written as:

$$y_{t,t+h} = \mathcal{X}'_{tj} \boldsymbol{\beta}_{j,h}^{\text{EN}} + \epsilon_{t+h}$$

where $\mathcal{X}_t = (1, \mathcal{X}_{1t}, \dots, \mathcal{X}_{Kt})'$ include the independent variable observations ($\mathbb{F}_t[y_{t,t+h}], \mathcal{Z}_{j,t}$) into a vector with “1” and $\boldsymbol{\beta}_{j,h}^{\text{EN}} = (\alpha_{j,h}, \beta_{j,h\mathbb{F}}, \text{vec}(\mathbf{B}_{j,h\mathcal{Z}}))' \equiv (\beta_0, \beta_1, \dots, \beta_K)'$ collects all the coefficients.

It is customary to standardize the elements of \mathcal{X}_t such that sample means are zero and sample standard deviations are unity. The coefficient estimates are then put back in their original scale by multiplying the slope coefficients by their respective standard deviations, and adding back the mean (scaled by slope coefficient over standard deviation.) The EN estimator incorporates both an L_1 and L_2 penalty:

$$\hat{\boldsymbol{\beta}}_{j,h}^{\text{EN}} = \underset{\beta_0, \beta_1, \dots, \beta_K}{\operatorname{argmin}} \left\{ \frac{1}{T_E} \sum_{\tau=1}^{T_E} \left(y_{\tau+h} - \mathcal{X}'_{\tau} \boldsymbol{\beta}_{j,h} \right)^2 + \underbrace{\lambda_1 \sum_{k=1}^K |\beta_{j,h,k}|}_{\text{LASSO}} + \underbrace{\lambda_2 \sum_{k=1}^K (\beta_{j,h,k})^2}_{\text{ridge}} \right\}$$

By minimizing the MSE over the training samples, we choose the optimal λ_1 and λ_2 values simultaneously.

In the implementation, the EN estimator is sometimes used as an input into the algorithm using the LSTM estimator. Specifically, we ensure that the machine forecast can only differ from the relevant benchmark if it demonstrably improves the pseudo out-of-sample prediction in the training samples *prior* to making a true out-of-sample forecast. Otherwise, the machine is replaced by the benchmark calculated up to time t . In some cases the benchmark is a survey forecast, in others it could be a historical mean value for the variable. However, for the implementation using LSTM, we also use the EN forecast as a benchmark.

E.2.2 Long Short-Term Memory (LSTM) Network

An LSTM network is a type of Recurrent Neural Network (RNN), which are neural networks used to learn about sequential data such as time series or natural language. In particular, LSTM

networks can learn long-term dependencies between across time periods by introducing hidden layers and memory cells to control the flow of information over longer time periods. The general case of the LSTM network with up to N hidden layers is defined as

$$\begin{aligned}
\underbrace{G^{\text{LSTM}}(\mathcal{X}_t, \beta_{j,h}^{\text{LSTM}})}_{1 \times 1} &= \underbrace{W^{(yh^N)}}_{1 \times D_{hN}} \underbrace{h_t^N}_{D_{hN} \times 1} + \underbrace{b_y}_{1 \times 1} & (\text{Output layer}) \\
\underbrace{h_t^n}_{D_{h^n} \times 1} &= \underbrace{o_t^n}_{D_{h^n} \times 1} \odot \tanh(\underbrace{c_t^n}_{D_{h^n} \times 1}) & (\text{Hidden layer}) \\
\underbrace{c_t^n}_{D_{h^n} \times 1} &= \underbrace{f_t^n}_{D_{h^n} \times 1} \odot \underbrace{c_{t-1}^n}_{D_{h^n} \times 1} + \underbrace{i_t^n}_{D_{h^n} \times 1} \odot \underbrace{\tilde{c}_t^n}_{D_{h^n} \times 1} & (\text{Final memory}) \\
\underbrace{\tilde{c}_t^n}_{D_{h^n} \times 1} &= \tanh(\underbrace{W^{(c^n h^{n-1})}}_{D_{h^n} \times D_{h^{n-1}}} \underbrace{h_t^{n-1}}_{D_{h^{n-1}} \times 1} + \underbrace{W^{(c^n h^n)}}_{D_{h^n} \times D_{h^n}} \underbrace{h_{t-1}^n}_{D_{h^n} \times 1} + \underbrace{b_{c^n}}_{D_{h^n} \times 1}) & (\text{New memory}) \\
\underbrace{f_t^n}_{D_{h^n} \times 1} &= \sigma(\underbrace{W^{(f^n h^{n-1})}}_{D_{h^n} \times D_{h^{n-1}}} \underbrace{h_t^{n-1}}_{D_{h^{n-1}} \times 1} + \underbrace{W^{(f^n h^n)}}_{D_{h^n} \times D_{h^n}} \underbrace{h_{t-1}^n}_{D_{h^n} \times 1} + \underbrace{b_{f^n}}_{D_{h^n} \times 1}) & (\text{Forget gate}) \\
\underbrace{i_t^n}_{D_{h^n} \times 1} &= \sigma(\underbrace{W^{(i^n h^{n-1})}}_{D_{h^n} \times D_{h^{n-1}}} \underbrace{h_t^{n-1}}_{D_{h^{n-1}} \times 1} + \underbrace{W^{(i^n h^n)}}_{D_{h^n} \times D_{h^n}} \underbrace{h_{t-1}^n}_{D_{h^n} \times 1} + \underbrace{b_{i^n}}_{D_{h^n} \times 1}) & (\text{Input gate}) \\
\underbrace{o_t^n}_{D_{h^n} \times 1} &= \sigma(\underbrace{W^{(o^n h^{n-1})}}_{D_{h^n} \times D_{h^{n-1}}} \underbrace{h_t^{n-1}}_{D_{h^{n-1}} \times 1} + \underbrace{W^{(o^n h^n)}}_{D_{h^n} \times D_{h^n}} \underbrace{h_{t-1}^n}_{D_{h^n} \times 1} + \underbrace{b_{o^n}}_{D_{h^n} \times 1}) & (\text{Output gate})
\end{aligned}$$

where $n = 1, \dots, N$ indexes each hidden layer. $h_t^n \in \mathbb{R}^{D_{h^n}}$ is the n -th *hidden layer*, where D_{h^n} is the number of *neurons* or *nodes* in the hidden layer. The 0-th layer is defined as the input data: $h_t^0 \equiv \mathcal{X}_t$. The memory cell c_t^n allows the LSTM network to retain information over longer time periods. The output gate o_t^n controls the extent to which the memory cell c_t^n maps to the hidden layer h_t^n . The forget gate f_t^n controls the flow of information carried over from the final memory in the previous timestep c_{t-1}^n . The input gate i_t^n controls the flow of information from the new memory cell \tilde{c}_t^n . The initial states for the hidden layers $(h_0^n)_{n=1}^N$ and memory cells $(c_0^n)_{n=1}^N$ are set to zeros.

$\sigma(\cdot)$ and $\tanh(\cdot)$ are *activation functions* that introduce non-linearities in the LSTM network, applied elementwise. $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ is the sigmoid function: $\sigma(x) = (1 + e^{-x})^{-1}$. $\tanh : \mathbb{R} \rightarrow \mathbb{R}$ is the hyperbolic tangent function: $\tanh(x) = \frac{e^{2x} - 1}{e^{2x} + 1}$. The \odot operator refers to elementwise multiplication.

$\beta_{j,h}^{\text{LSTM}} \equiv (((\text{vec}(W^{(g^n h^{n-1})})', \text{vec}(W^{(g^n h^n)})', b_{g^n}')_{g \in \{c, f, i, o\}})_{n=1}^N, \text{vec}(W^{(yh^N)})', b_y)'$ are parameters to be estimated. We will refer to parameters indexed with W as *weights*; parameters indexed with b are *biases*. We estimate the parameters $\beta_{j,h}^{\text{LSTM}}$ for the LSTM network using Stochastic Gradient Decent (SGD), which is an iterative algorithm for minimizing the loss function and proceeds as follows:

1. *Initialization.* Fix a random seed R and draw a starting value of the parameters $\beta_{j,h}^{(0)}$ randomly, where the superscript (0) in parentheses indexes the iteration for an estimate of $\beta_{j,h}^{\text{LSTM}}$.

- (a) Initialize the input weights $W^{(g^n h^{n-1})} \in \mathbb{R}^{D_{h^n} \times D_{h^{n-1}}}$ for $g \in \{c, f, i, o\}$ using the *Glorot* initializer. Draw randomly from a uniform distribution with zero mean and a variance that depends on the dimensions of the matrix:

$$W_{ij}^{(g^n h^{n-1})} \stackrel{iid}{\sim} U \left[-\sqrt{\frac{6}{D_{h^n} + D_{h^{n-1}}}}, \sqrt{\frac{6}{D_{h^n} + D_{h^{n-1}}}} \right]$$

for each $i = 1, \dots, D_{h^n}$ and $j = 1, \dots, D_{h^{n-1}}$.

- (b) Initialize the recurrent weights $W^{(g^n h^n)} \in \mathbb{R}^{D_{h^n} \times D_{h^n}}$ for $g \in \{c, f, i, o\}$ using the *Orthogonal* initializer. Use the orthogonal matrix obtained from the QR decomposition of a $D_{h^n} \times D_{h^n}$ matrix of random numbers drawn from a standard normal distribution.
- (c) Initialize biases $(b_{g^n})_{g \in \{c, f, i, o\}}$, hidden layers h_0^n , and memory cells c_0^n with zeros.

2. *Mini-batches.* Prepare the input data by dividing the training sample into a collection of *mini-batches*.

- (a) Suppose that we have a multi-variate time-series training sample with dimensions (T_E, K) whose time steps t are indexed by $t = 1, \dots, T_E$ and K is the number of predictors. We transform this training sample into a 3-D tensor with dimensions (N_S, M, K) where

- N_S = Total number of sequences in training sample
- M = Sequence length, i.e., number of time steps in each sequence
- K = Input size, i.e., number of predictors in each time step

This can be done by creating overlapping sequences from the time series:

- Sequence 1 contains time steps $1, \dots, M$
- Sequence 2 contains time steps $2, \dots, M + 1$
- Sequence 3 contains time steps $3, \dots, M + 2$
- ...
- Sequence $T_E - M$ contains time steps $T_E - M, \dots, T_E - 1$
- Sequence $N_S = T_E - M + 1$ contains time steps $T_E - M + 1, \dots, T_E$

- (b) Randomly shuffle the N_S sequences by randomly sampling a permutation of the sequences without replacement.
- (c) Partition the N_S shuffled sequences into $\lceil N_S/N_B \rceil$ mini-batches. We partition the N_S sequences in the training sample $((N_S, M, K)$ tensor) into a list of $\lceil N_S/N_B \rceil$ mini-batches. A mini-batch is a (N_B, M, K) -dimensional tensor containing N_B out of N_S randomly shuffled sequences.⁸ Let $B^{(1)}, \dots, B^{[N_S/N_B]}$ denote the list of mini-batches.

⁸When N_S/N_B is not a whole number, $\lceil N_S/N_B \rceil$ of the mini-batches will be 3-D tensors with dimensions (N_B, M, K) . One batch will contain leftover sequences and will have dimensions $(N_S \% N_B, M, K)$ where $\%$ is the modulus operator.

- N_S = Total number of sequences in training sample
- N_B = Mini-batch size, i.e., number of sequences in each partition.
- M = Sequence length, i.e., number of time steps in each sequence
- K = Input size, i.e., number of predictors in each time step

3. Repeat until the stopping condition is satisfied ($k = 1, 2, 3, \dots$):

- (a) *Dropout*. Apply dropout to the mini-batch. To obtain the n -th hidden layer under dropout, multiply the current value of the $n - 1$ -th hidden layer h_t^{n-1} and the lagged value of the n -th hidden layer h_{t-1}^n with binary masks $r_{t,h_t^{n-1}}^{(k)} \in \mathbb{R}^{D_{h^{n-1}}}$ and $r_{t,h_{t-1}^n}^{(k)} \in \mathbb{R}^{D_{h^n}}$, respectively:

$$\begin{aligned} \underbrace{\bar{h}_t^{n-1}}_{D_{h^{n-1}} \times 1} &= \underbrace{r_{t,h_t^{n-1}}^{(k)}}_{D_{h^{n-1}} \times 1} \odot \underbrace{h_t^{n-1}}_{D_{h^{n-1}} \times 1}, \quad r_{t,h_t^{n-1},i}^{(k)} \stackrel{iid}{\sim} \text{Bernoulli}(p_{h_t^{n-1}}), \quad i = 1, \dots, D_{h^{n-1}} \\ \underbrace{\bar{h}_{t-1}^n}_{D_{h^n} \times 1} &= \underbrace{r_{t,h_{t-1}^n}^{(k)}}_{D_{h^n} \times 1} \odot \underbrace{h_{t-1}^n}_{D_{h^n} \times 1}, \quad r_{t,h_{t-1}^n,i}^{(k)} \stackrel{iid}{\sim} \text{Bernoulli}(p_{h_{t-1}^n}), \quad i = 1, \dots, D_{h^n} \end{aligned}$$

where $t \in B^{(k)}$ and $n = 1, \dots, N$ indexes the hidden layer and it is understood that the 0-th layer is the input vector $h_t^0 \equiv \mathcal{X}_t$. $p_{h_t^{n-1}}, p_{h_{t-1}^n} \in [0, 1]$ is the probability that time t nodes in the $n - 1$ -th hidden layer and time $t - 1$ nodes in the n -th hidden layer are retained, respectively.

- (b) *Stochastic Gradient*. Average the gradient over observations in the mini-batch

$$\nabla L(\beta_{j,h}^{(k-1)}, \mathbf{X}_{B^{(k)}}, \boldsymbol{\lambda}^{\text{LSTM}}) = \frac{1}{M} \sum_{t \in B^{(k)}} \nabla L(\beta_{j,h}^{(k-1)}, \mathbf{X}_t, \boldsymbol{\lambda}^{\text{LSTM}})$$

where $\nabla L(\beta_{j,h}^{(k-1)}, \mathbf{X}_t, \boldsymbol{\lambda}^{\text{LSTM}})$ is the gradient of the loss function with respect to the parameters $\beta_{j,h}^{(k-1)}$, evaluated at the time t observation $\mathbf{X}_t = (y_{t,t+h}, \hat{\mathcal{X}}_t)'$ after applying dropout.

- (c) *Learning rate shrinkage*. Update the parameters to $\beta_{j,h}^{(k)}$ using the Adaptive Moment Estimation (Adam) algorithm. The method uses the first and second moments of the gradients to shrink the overall learning rate to zero as the gradient approaches zero.

$$\beta_{j,h}^{(k)} = \beta_{j,h}^{(k-1)} - \gamma \frac{m^{(k)}}{\sqrt{v^{(k)}} + \varepsilon}$$

where $m^{(k)}$ and $v^{(k)}$ are weighted averages of first two moments of past gradients:

$$\begin{aligned} m^{(k)} &= \frac{1}{1 - \pi_1^k} (\pi_1 m^{(k-1)} + (1 - \pi_1) \nabla L(\beta_{j,h}^{(k-1)}, \mathbf{X}_{B^{(k)}}, \boldsymbol{\lambda}^{\text{LSTM}})) \\ v^{(k)} &= \frac{1}{1 - \pi_2^k} (\pi_2 v^{(k-1)} + (1 - \pi_2) \nabla L(\beta_{j,h}^{(k-1)}, \mathbf{X}_{B^{(k)}}, \boldsymbol{\lambda}^{\text{LSTM}})^2) \end{aligned}$$

π^k denotes the k -th power of $\pi \in (0, 1)$, and $/$, $\sqrt{\cdot}$, and $(\cdot)^2$ are applied element-wise. The default values of the hyperparameters are $m^{(0)} = v^{(0)} = 0$ (initial moment vectors), $\gamma = 0.001$ (initial learning rate), $(\pi_1, \pi_2) = (0.9, 0.999)$ (decay rates), and $\varepsilon = 10^{-7}$ (prevent zero denominators).

(d) *Stopping Criteria.* Stop iterating and return $\beta_{j,h}^{(k)}$ if one of the following holds:

- *Early stopping.* At each iteration, use the updated $\beta_{j,h}^{(k)}$ to calculate the loss from the validation sample. Stop when the validation loss has not improved for S steps, where S is a “patience” hyperparameter. By updating the parameters for fewer iterations, early stopping shrinks the final parameters $\beta_{j,h}$ towards the initial guess $\beta_{j,h}^{(0)}$, and at a lower computational cost than ℓ_2 regularization.
- *Maximum number of epochs.* Stop if the number of iterations reaches the maximum number of epochs E . An epoch happens when the full set of the training sample has been used to update the parameters. If the training sample has T_E observations and each mini-batch has M observations, then each epoch would contain $\lceil T_E/M \rceil$ iterations (after rounding up as needed). So the maximum number of iterations is bounded by $E \times \lceil T_E/M \rceil$.

4. *Ensemble forecasts.* Repeat steps 1. and 2. over different random seeds R and save each of the estimated parameters $\hat{\beta}_{j,h,T_E}^{LSTM}(\mathbf{X}_{T_E}, \boldsymbol{\lambda}^{LSTM}, R)$. Then construct out-of-sample forecasts using the top 10 out of 20 starting values with the best performance in the validation sample. Ensemble can be considered as a regularization method because it aims to guard against overfitting by shrinking the forecasts toward the average across different random seeds. The random seed affects the random draws of the parameter’s initial starting value $\beta_{j,h}^{(0)}$, the sequences selected in each mini-batch $B^{(k)}$, and the dropout mask $r_t^{(k)}$.

E.2.3 Hyperparameters

Let $\boldsymbol{\lambda}^{LSTM} \equiv [\lambda_1, \lambda_2, \gamma, \pi_1, \pi_2, p, N, (D_{h^n})_{n=1}^N, M, E, S]'$ collect all the hyperparameters that control the LSTM network’s complexity and prevent the model from overfitting the data. The number of hidden layers N and the number of neurons D_{h^1}, \dots, D_{h^N} in each hidden layer are hyperparameters that characterize the network’s *architecture*. The hyperparameters are estimated by minimizing the mean-square loss over pseudo out-of-sample forecast errors generated from rolling regressions through the validation sample.

Table A.18 reports the hyperparameters for the LSTM network and its estimation. Hyperparameters reported as a range or a set of values are cross-validated. We ask the machine to cross validate the L_1 and L_2 penalties by adjusting the relative weight on the L_1 penalty, α , and the overall multiplier, λ , within the range of values listed in the table. α and λ are converted back to the L_1 and L_2 penalties using the mapping $\lambda_1 = \alpha\lambda$ and $\lambda_2 = \frac{1-\alpha}{2}\lambda$. The window lengths for the training sample T_E and validation sample T_V denote the number of quarterly observations. $T_V = \textit{Expand}$ denotes the case of recursively expanding windows, where pseudo out-of-sample

forecasts are generated from rolling estimates based on a training sample that expands recursively over time starting from 1984Q1 for returns and price growth, 1995Q1 for earnings growth.

Table A.18: Hyperparameters for the Machine Algorithm

(1) Machine Forecast: Stock Returns and Price Growth (1 Year Ahead)

Description / Horizon	Returns (1 Year)	Price Growth (1 Year)
(a) Elastic Net (λ^{EN})		
Weight on L_1 penalty (α)	[0.01, 0.99]	[0.01, 0.99]
Multiplier on L_1 and L_2 penalties (λ)	$[10^{-2}, 10]$	$[10^{-2}, 10]$
Size of training sample (T_E)	4, 5, 6, 7	4, 5, 6, 7
Size of validation sample (T_V)	4, 5, 6, 7, <i>Expand</i>	4, 5, 6, 7, <i>Expand</i>
(b) Long Short-Term Memory Network (λ^{LSTM})		
Weight on L_1 penalty (α)	[0.01, 0.99]	[0.01, 0.99]
Multiplier on L_1 and L_2 penalties (λ)	$[10^{-4}, 10^{-2}]$	$[10^{-4}, 10^{-2}]$
Initial learning rate (γ)	0.001	0.001
Gradient decay rates (π_1, π_2)	0.9, 0.999	0.9, 0.999
Probability of dropout ($p_{\mathcal{X}}, p_{h^n}$)	0.8, 0.5	0.8, 0.5
Number of hidden layers (N)	1	1
Number of neurons ($(D_{h^n})_{n=1}^N$)	4	4
Mini-batch size (N_B)	4	4
Sequence length (M)	4	4
Patience for early stopping (S)	10	20
Maximum number of epochs (E)	10,000	10,000
Random seeds (R)	1, 2, ..., 20	1, 2, ..., 20
Size of training sample (T_E)	5, 7	5, 7
Size of validation sample (T_V)	5, 7, 20	5, 7, 20

(2) Machine Forecast: Stock Returns (2, 3, 4, and 5 Years Ahead)

Description / Horizon	2 Years	3 Years	4 Years	5 Years
(a) Elastic Net (λ^{EN})				
Weight on L_1 penalty (α)	[0.01, 0.99]	[0.1, 0.9]	[0.1, 0.9]	[0.1, 0.9]
Multiplier on L_1 and L_2 (λ)	$[10^{-2}, 10]$	$[10^{-1}, 10]$	$[10^{-1}, 10]$	$[10^{-1}, 10]$
Size of training sample (T_E)	4, 5, 6, 7	4, 5, 6, 7	4, 5, 6, 7	4, 5, 6, 7
Size of validation sample (T_V)	4, 5, 6, 7 <i>Expand</i>	4, 5, 6, 7 <i>Expand</i>	4, 5, 6, 7 <i>Expand</i>	4, 5, 6, 7 <i>Expand</i>
(b) Long Short-Term Memory Network (λ^{LSTM})				
Weight on L_1 penalty (α)	[0.01, 0.99]	[0.01, 0.99]	[0.01, 0.99]	[0.01, 0.99]
Multiplier on L_1 and L_2 (λ)	$[10^{-4}, 10^{-2}]$	$[10^{-4}, 10^{-2}]$	$[10^{-4}, 10^{-2}]$	$[10^{-4}, 10^{-2}]$
Initial learning rate (γ)	0.001	0.001	0.001	0.001
Gradient decay rates (π_1, π_2)	0.9, 0.999	0.9, 0.999	0.9, 0.999	0.9, 0.999
Probability of dropout ($1 - p$)	0.5	0.5	0.5	0.05
Number of hidden layers (N)	1	1	1	1
Number of neurons ($(D_{h^n})_{n=1}^N$)	4	4	4	4
Mini-batch size (M)	2	2	2	2
Patience for early stopping (S)	3	3	3	20
Maximum number of epochs (E)	1,000	1,000	1,000	10,000
Random seeds (R)	1, 2, ..., 20	1, 2, ..., 20	1, 2, ..., 20	1, 2, ..., 20
Size of training sample (T_E)	3, 5, 7	3, 5, 7	3, 5, 7	3, 5, 7
Size of validation sample (T_V)	3, 5, 7	3, 5, 7	3, 5, 7	3, 5, 7

Notes: This table reports the hyperparameters considered in the machine learning algorithm for each estimator.

(3) Machine Forecast: Earnings Growth (1, 2, 3, and 4 Years Ahead)

Description / Horizon	1 Year	2 Years	3 Years	4 Years
(a) Elastic Net (λ^{EN})				
Weight on L_1 penalty (α)	[0.1, 0.9]	[0.1, 0.9]	[0.1, 0.9]	[0.1, 0.9]
Multiplier on L_1 and L_2 (λ)	[1, 10]	[1, 10]	[1, 10]	[1, 10]
Size of training sample (T_E)	4, 6, 8, 10, 12	4, 6, 8, 10, 12	4, 6, 8, 10, 12	4, 6, 8, 10, 12
Size of validation sample (T_V)	4, 6, 8, 10, 12 <i>Expand</i>	4, 6, 8, 10, 12 <i>Expand</i>	4, 6, 8, 10, 12 <i>Expand</i>	4, 6, 8, 10, 12 <i>Expand</i>
(b) Long Short-Term Memory Network (λ^{LSTM})				
Weight on L_1 penalty (α)	[0.01, 0.99]	[0.01, 0.99]	[0.01, 0.99]	[0.01, 0.99]
Multiplier on L_1 and L_2 (λ)	$[10^{-4}, 10^{-2}]$	$[10^{-4}, 10^{-2}]$	$[10^{-4}, 10^{-2}]$	$[10^{-4}, 10^{-2}]$
Initial learning rate (γ)	0.001	0.001	0.001	0.001
Gradient decay rates (π_1, π_2)	0.9, 0.999	0.9, 0.999	0.9, 0.999	0.9, 0.999
Probability of dropout ($p_{\mathcal{X}}, p_{h^n}$)	0.05, 0.05	0.05, 0.05	0.05, 0.05	0.05, 0.05
Number of hidden layers (N)	1	1	1	1
Number of neurons ($D_{h^n}_{n=1}$)	4	4	4	4
Mini-batch size (N_B)	4	4	4	4
Sequence length (M)	4	4	4	4
Patience for early stopping (S)	20	20	20	20
Maximum number of epochs (E)	10,000	10,000	10,000	10,000
Random seeds (R)	1, 2, ..., 20	1, 2, ..., 20	1, 2, ..., 20	1, 2, ..., 20
Size of training sample (T_E)	4, 8, 12	4, 8, 12	4, 8, 12	4, 8, 12
Size of validation sample (T_V)	4, 8, 12, 20	4, 8, 12, 20	4, 8, 12, 20	4, 8, 12, 20

(4) Machine Forecast: Long-Term Growth (LTG)

Description / Horizon	4-to-5-years	0-to-5-years	1-to-10-years
(a) Elastic Net (λ^{EN})			
Weight on L_1 penalty (α)	[0.1, 0.9]	[0.1, 0.9]	[0.1, 0.9]
Multiplier on L_1 and L_2 (λ)	[1, 10]	[1, 10]	[1, 10]
Size of training sample (T_E)	6, 8, 10, 12	6, 8, 10, 12	6, 8, 10, 12
Size of validation sample (T_V)	6, 8, 10, 12 <i>Expand</i>	6, 8, 10, 12 <i>Expand</i>	6, 8, 10, 12 <i>Expand</i>
(b) Long Short-Term Memory Network (λ^{LSTM})			
Weight on L_1 penalty (α)	[0.01, 0.99]	[0.01, 0.99]	[0.01, 0.99]
Multiplier on L_1 and L_2 (λ)	$[10^{-5}, 10^{-1}]$	$[10^{-4}, 10^{-2}]$	$[10^{-5}, 10^{-1}]$
Initial learning rate (γ)	0.001	0.001	0.001
Gradient decay rates (π_1, π_2)	0.9, 0.999	0.9, 0.999	0.9, 0.999
Probability of dropout ($p_{\mathcal{X}}, p_{h^n}$)	0.05, 0.05	0.05, 0.05	0.05, 0.05
Number of hidden layers (N)	1	1	1
Number of neurons ($D_{h^n}_{n=1}$)	4	4	4
Mini-batch size (N_B)	4	4	4
Sequence length (M)	4	4	4
Patience for early stopping (S)	80	20	20
Maximum number of epochs (E)	10,000	10,000	10,000
Random seeds (R)	1, 2, ..., 20	1, 2, ..., 20	1, 2, ..., 20
Size of training sample (T_E)	3, 7, 12; $T_E = 3$ if $T_V = 20$	4, 8, 12	3, 7, 12; $T_E = 3$ if $T_V = 20$
Size of validation sample (T_V)	3, 7, 12, 20	4, 8, 12, 20	3, 7, 12, 20

Notes: This table reports the hyperparameters considered in the machine learning algorithm for each estimator. The LSTM forecast is replaced with the corresponding EN forecast if the LSTM cannot improve on the EN over the pseudo out-of-sample predictions. The LSTM forecast for 0-to-5 year ahead LTG is implemented as an estimation using forecast errors as the dependent variable.

(5) Machine Forecast: Inflation and Employment Growth (1 Year Ahead)

Description / Horizon	Inflation (1 Year)	Employment Growth (1 Year)
(a) Elastic Net (λ^{EN})		
Weight on L_1 penalty (α)	[0.01, 0.99]	[0.01, 0.99]
Multiplier on L_1 and L_2 penalties (λ)	$[10^{-2}, 3]$	$[10^{-1}, 10]$
Size of training sample (T_E)	3, 4, 5, 6, 7	3, 4, 5, 6, 7
Size of validation sample (T_V)	6, 7, ..., 14, 15	6, 7, ..., 14, 15
(b) Long Short-Term Memory Network (λ^{LSTM})		
Weight on L_1 penalty (α)	[0.01, 0.99]	[0.01, 0.99]
Multiplier on L_1 and L_2 penalties (λ)	$[10^{-4}, 10^{-2}]$	$[10^{-4}, 10^{-2}]$
Initial learning rate (γ)	0.001	0.001
Gradient decay rates (π_1, π_2)	0.9, 0.999	0.9, 0.999
Probability of dropout ($p_{\mathcal{X}}, p_{h^n}$)	0.8, 0.5	0.005, 0.005
Number of hidden layers (N)	1	1
Number of neurons (D_{h^n}) $_{n=1}^N$	4	4
Mini-batch size (N_B)	4	4
Sequence length (M)	4	4
Patience for early stopping (S)	20	5
Maximum number of epochs (E)	10,000	10,000
Random seeds (R)	1, 2, ..., 20	1, 2, ..., 20
Size of training sample (T_E)	5, 7	3, 5
Size of validation sample (T_V)	6, 9, 12, 15	6, 9, 12

Notes: This table reports the hyperparameters considered in the machine learning algorithm for each estimator.

E.2.4 Machine Variables to Be forecast

Returns and price growth When evaluating the MSE ratio of the machine relative to that of a benchmark survey, we use the machine forecast for the return or price growth measure that most closely corresponds to the concept that survey respondents are asked to predict:

1. CFO survey asks respondents about their expectations for the S&P 500 return over the next 12 months. Following Nagel and Xu (2021), we interpret the survey to be asking about the one-year CRSP value-weighted return (including dividends) from the current survey month to the same month one year ahead.
2. Gallup/UBS survey respondents report the return (including dividends) they expect on their own portfolio one year ahead. We interpret the survey to be asking about the one-year CRSP value-weighted return (including dividends) from the current survey month to the same month one year ahead.
3. Livingston survey respondents provide 12-month ahead forecasts of the S&P 500 index. We convert the level forecast to price growth forecast by taking the log difference between the 12-month ahead level forecast and the nowcast of the S&P 500 index for the current survey month. Therefore, we interpret the survey to be asking about the one-year price growth in the S&P 500 index.
4. Bloomberg Consensus Forecasts asks survey respondents about the end-of-year closing value of the S&P 500 index. We interpret the survey to be asking about the h -month price growth in the S&P 500 index. The horizon of the forecast changes depending on when in the year the panelists are answering the survey.
5. Michigan Survey of Consumers (SOC) asks respondents about their perceived probability that an investment in a diversified stock fund would increase in value in the year ahead. We interpret the question to be asking about the one-year price growth in the S&P 500 index.
6. Conference Board (CB) survey asks respondents about their categorical belief on whether they expect stock prices to increase, decrease, or stay the same over the next year. We interpret the question to be asking about the one-year price growth in the S&P 500 index.
7. When the machine forecast is compared with the historical mean as the benchmark, we compare the forecasts for annualized CRSP return. For forecast horizons longer than one year, we compare the machine forecast of annualized cumulative log CRSP return from the time of the forecast to the end of the forecast horizon less the current short rate against the historical mean of the same variable.

Earning Growth For earning growth forecasts, we use a quarterly S&P 500 total earnings series that combines data from S&P Global, Shiller, and the S&P 500 divisor, as described above. Since the machine learning algorithm has been adapted to a quarterly forecasting frequency, we use the quarterly series before the monthly interpolation. The quarterly series spans the period 1959:Q1 to 2021:Q4.

For Long-Term Growth (LTG) forecasts, IBES defines LTG as the “expected annual increase in operating earnings over the company’s next full business cycle. These forecasts refer to a period of between three to five years.” We compare survey responses of LTG against machine forecasts under alternative interpretations of LTG. First, we consider machine forecasts of annual 5-year forward growth, i.e., annual earnings growth from four to five years ahead. Second, we consider machine forecasts of annualized 5-year growth, i.e., annual earnings growth from current quarter to five years ahead, following the interpretation in Bordalo et al. (2019). Third, we consider machine forecasts of annualized earnings growth from one to 10 years ahead, following the interpretation in Nagel and Xu (2021)

Inflation We construct forecasts of annual inflation defined as

$$\pi_{t+4,t} = \log \left(\frac{PGDP_{t+4}}{PGDP_t} \right)$$

where $PGDP_t$ is the quarterly level of the chain-weighted GDP price index. Following Coibion and Gorodnichenko (2015), we use the vintage of data that are available four quarters after the period being forecast.

Employment growth We construct forecasts of annual employment growth defined as

$$y_{t+4,t} = \log \left(\frac{EMP_{t+4}}{EMP_t} \right)$$

Following Coibion and Gorodnichenko (2015), we use the vintage of data that are available four quarters after the period being forecast.

E.2.5 Machine Input Data: Predictor Variables

The vector $\mathbf{Z}_{jt} \equiv (y_t, \hat{\mathbf{G}}'_t, \mathbf{W}'_{jt})'$ is an $r = 1 + r_G + r_W$ vector which collects the data at time t with $\mathcal{Z}_{jt} \equiv (y_t, \dots, y_{t-p_y}, \hat{\mathbf{G}}'_t, \dots, \hat{\mathbf{G}}'_{t-p_G}, \mathbf{W}'_{jt}, \dots, \mathbf{W}'_{jt-p_W})'$ a vector of contemporaneous and lagged values of \mathbf{Z}_{jt} , where p_y, p_G, p_W denote the total number of lags of $y_t, \hat{\mathbf{G}}'_t, \mathbf{W}'_{jt}$, respectively. The predictors below are listed as elements of $y_t, \hat{\mathbf{G}}'_{jt}$, or \mathbf{W}'_{jt} for variables.

Stock return and price growth predictor variables and specifications For y_j equal to CRSP value-weighted returns or S&P 500 price index growth, we first predict the one-year log stock return or price growth that is expected to occur h quarters into the future from time

$t + h - 4$ to $t + h$, i.e., $\mathbb{E}_t[r_{t+h-4,t+h}]$. For horizons longer than one year, since the h -quarter long horizon return is the sum of one-year returns between time t to $t + h$, we first forecast the forward one-year returns separately and then add the components together to get machine forecasts of h -quarter long horizon returns. The forecasting model considers the following variables:

1. $\mathbf{G}_{M,t-k}$, for $k = 0, 1$ are factors formed from a real-time macro dataset \mathcal{D}^M with 92 real-time macro series; includes both monthly and quarterly series, with monthly series converted to quarterly according to the method described in the data appendix.
2. $\mathbf{G}_{F,t-k}$, for $k = 0, 1$ are factors formed from a financial data set \mathcal{D}^F with 147 monthly financial series.
3. $\mathbf{G}_{D,t-k}^Q$, for $k = 0$ are quarterly factors formed from a daily financial dataset \mathcal{D}^D of 87 daily financial indicators. The raw daily series are first converted to daily factors $\mathbf{G}_{D,t}(\mathbf{w})$ and the daily factors are aggregated up to quarterly observations $\mathbf{G}_{D,t}^Q(\mathbf{w})$ using a weighted average of daily factors, with the weights \mathbf{w} dependent on two free parameters that are chosen to minimize the sum of squared residuals in a regression of $y_{t,t+h}$ on $\mathbf{G}_{D,t}(\mathbf{w})$.
4. *LDA topics* $F_{k,t}$, for topic $k = 1, 2, \dots, 50$. The value of the topic k at time t is the average weights of topic k of all articles published at t .
5. *Macro data surprises* from the money market survey. The macro news include, GDP growth (Q/Q percentage change), core CPI (Month/Month change), unemployment rate (percentage point), and nonfarm payroll (month/month change). We include first release, second release, and final release for GDP growth. This constitutes six macro data surprises per quarter.
6. *FOMC surprises* are defined as the changes in the current-month, 1, 2, 3, 4, 6, 12, and 24 month-ahead federal funds futures (FFF) contract rate and the changes in the 1, 2, 3, 4, and 8 quarter-ahead Eurodollar (ED) futures contracts, from 10 minutes before to 20 minutes after each FOMC announcement. When benchmarking against a survey, we use the last FOMC meeting before the survey deadline to compute surprises. For surveys that do not have a clear deadline, we compute surprises using from the last FOMC in the first month of the quarter. When benchmarking against moving average, we use the last FOMC meeting before the end of the first month in each quarter to compute surprises. This leaves 10 FOMC surprise variables per quarter.
7. *Stock market jumps* are accumulated 30-minute window negative and positive jumps in the S&P 500 around news events over the previous quarter.
8. y_{t-k} for $k = 1, 2$ are lags of the dependent variable. For price growth forecasts, include lags of CRSP returns. We include this term for forecast horizons of 1 and 2 years.

9. $\bar{\mu}_{t-k}$ for $k = 1, 2$ is the historical mean of the dependent variable calculated up to the middle month of the quarter at time t . The initial period is 1959Q1. For price growth forecasts, include the historical mean of CRSP returns.
10. *Long-term growth of earnings*: Annualized 5-year log growth rate of quarterly S&P 500 total earnings at the end of the previous quarter. We include this term for forecast horizons of 1 and 2 years.
11. *Short rates*. When forecasting returns or price growth, the machine controls for the current nominal short rate, $\log(1 + 3MTB_t/100)$, imposing a unit coefficient. This is equivalent to forecasting the future return minus the current short rate.
12. For LSTM forecasts only: EN forecast of the variable to be predicted. Missing EN forecasts during earlier periods are imputed with the historical mean of the dependent variable. For price growth forecasts, include the EN forecast of CRSP returns.

The 92 macro series in \mathcal{D}^M are selected to represent broad categories of macroeconomic time series. The majority of these are real activity measures: real output and income, employment and hours, consumer spending, housing starts, orders and unfilled orders, compensation and labor costs, and capacity utilization measures. The dataset also includes commodity and price indexes and a handful of bond and stock market indexes, and foreign exchange measures. The financial dataset \mathcal{D}^f is an updated monthly version of the of 147 variables comprised solely of financial market time series used in Ludvigson and Ng (2007). These data include valuation ratios such as the dividend-price ratio and earnings-price ratio, growth rates of aggregate dividends and prices, default and term spreads, yields on corporate bonds of different ratings grades, yields on Treasuries and yield spreads, and a broad cross-section of industry, size, book-market, and momentum portfolio equity returns.⁹ The 87 daily financial indicators in \mathcal{D}^D include daily time series on commodities spot prices and futures prices, aggregate stock market indexes, volatility indexes, credit spreads and yield spreads, and exchange rates.

After constructing machine forecasts of forward one-year stock returns $\mathbb{E}_t[r_{t+h-4,t+h}]$, we construct machine forecasts of cumulative h -quarter long horizon returns between time t to $t + h$, i.e., $\mathbb{E}_t[r_{t,t+h}]$, as the sum of the forward one-year return expectations up to h quarters ahead:

$$\mathbb{E}_t[r_{t,t+h}] = \mathbb{E}_t[r_{t,t+4}] + \mathbb{E}_t[r_{t+4,t+8}] + \cdots + \mathbb{E}_t[r_{t+h-4,t+h}]$$

Earning growth predictor variables and specifications For earning growth forecasts, we first detrend the (log) earnings level in real time by, starting with an initial sample, recursively running the following regression at each point in time t

$$\log(\text{earnings}_t) = \alpha_t + \beta_t t + y_t$$

⁹A detailed description of the series is given in the Data Appendix of the online supplementary file at www.sydneyludvigson.com/s/ucc_data_appendix.pdf

For y_t equal to the detrended (log) earning level, we construct a forecasted value for y_t , denoted $\hat{y}_{t|t-h}$, based on information known up to time t using the following variables:

1. $\mathbf{G}_{M,t-k}$, for $k = 0, 1$ are factors formed from a real-time macro dataset \mathcal{D}^M with 92 real-time macro series; includes both monthly and quarterly series, with monthly series converted to quarterly according to the method described in the data appendix.
2. $\mathbf{G}_{F,t-k}$, for $k = 0, 1$ are factors formed from a financial data set \mathcal{D}^F with 147 monthly financial series.
3. $\mathbf{G}_{D,t-k}^Q$, for $k = 0$ are quarterly factors formed from a daily financial dataset \mathcal{D}^D of 87 daily financial indicators. The raw daily series are first converted to daily factors $\mathbf{G}_{D,t}(\mathbf{w})$ and the daily factors are aggregated up to quarterly observations $\mathbf{G}_{D,t}^Q(\mathbf{w})$ using a weighted average of daily factors, with the weights \mathbf{w} dependent on two free parameters that are chosen to minimize the sum of squared residuals in a regression of y_t on $\mathbf{G}_{D,t}(\mathbf{w})$.
4. *LDA factors* $F_{k,t-j}$, for topic $k = 1, 2, \dots, 50$ and $j = 0, 1$. The value of the topic k at time t is the average weights of topic k of all articles published at t .
5. *Macro data surprises* from the money market survey. The macro news include, GDP growth (Q/Q percentage change), core CPI (Month/Month change), unemployment rate (percentage point), and nonfarm payroll (month/month change). We include first release, second release, and final release for GDP growth. This constitutes six macro data surprises per quarter.
6. *FOMC surprises* are defined as the changes in the current-month, 1, 2, 3, 4, 6, 12, and 24 month-ahead federal funds futures (FFF) contract rate and the changes in the 1, 2, 3, 4, and 8 quarter-ahead Eurodollar (ED) futures contracts, from 10 minutes before to 20 minutes after each FOMC announcement. When benchmarking against a survey, we use the last FOMC meeting before the survey deadline to compute surprises. For surveys that do not have a clear deadline, we compute surprises using from the last FOMC in the first month of the quarter. When benchmarking against moving average, we use the last FOMC meeting before the end of the first month in each quarter to compute surprises. This leaves 10 FOMC surprise variables per quarter.
7. *Stock market jumps* are accumulated 30-minute window negative and positive jumps in the S&P 500 around news events over the previous quarter.
8. y_{t-k} for $k = 1, 2$ are lags of the dependent variable.
9. $\bar{\mu}_{t-k}$ for $k = 1, 2$ is the historical mean of the dependent variable calculated up to time t . The initial period is 1959Q1.
10. $\mathbb{F}_{t-k}[y_{t+h-k}]$ for $k = 1, 2$ are lags of the survey forecasts of the dependent variable.

11. For LSTM forecasts only: EN forecast of the variable to be predicted. Missing EN forecasts during earlier periods are imputed with the historical mean of the dependent variable.

The 92 macro series in \mathcal{D}^M are selected to represent broad categories of macroeconomic time series. The majority of these are real activity measures: real output and income, employment and hours, consumer spending, housing starts, orders and unfilled orders, compensation and labor costs, and capacity utilization measures. The dataset also includes commodity and price indexes and a handful of bond and stock market indexes, and foreign exchange measures. The financial dataset \mathcal{D}^f is an updated monthly version of the of 147 variables comprised solely of financial market time series used in Ludvigson and Ng (2007). These data include valuation ratios such as the dividend-price ratio and earnings-price ratio, growth rates of aggregate dividends and prices, default and term spreads, yields on corporate bonds of different ratings grades, yields on Treasuries and yield spreads, and a broad cross-section of industry, size, book-market, and momentum portfolio equity returns.¹⁰ The 87 daily financial indicators in \mathcal{D}^D include daily time series on commodities spot prices and futures prices, aggregate stock market indexes, volatility indexes, credit spreads and yield spreads, and exchange rates.

After we obtain the machine forecast for the detrended level of earnings, y , we obtain the h -horizon machine earnings growth forecast (from $t - h$ to t denoted $\mathbb{E}_{t-h} [\Delta \log (\text{earnings}_t^M)]$) by constructing

$$\mathbb{E}_{t-h} [\Delta \log (\text{earnings}_t^M)] = \hat{\alpha}_{t-h} + \hat{\beta}_{t-h}t + \hat{y}_{t|t-h}^M - \log (\text{earnings}_{t-h})$$

where $\log (\text{earnings}_{t-h})$ is the realized log earning level at time $t - h$, and $\hat{y}_{t|t-h}^M$ is the machine forecast of the detrended log earnings based on information up to time $t - h$. Following De La O and Myers (2021) and Bordalo et al. (2019), we assume that the survey respondents and the machine could observe $\log (\text{earnings}_{t-h})$, the log earnings level for the quarter on which the forecast is made.

To use this approach to forecast the 20-quarter ahead annual forward earnings i.e., (from $t - 4$ to t on basis of information at $t - 20$), we would construct

$$\mathbb{E}_{t-20} [\log (\text{earnings}_t^M)] = \hat{\alpha}_{t-20} + \hat{\beta}_{t-20}t + \hat{y}_{t|t-20}^M.$$

To construct 20-quarter ahead annual earnings growth forecast we compute

$$\mathbb{E}_{t-20} [\log (\text{earnings}_{t-4}^M)] = \hat{\alpha}_{t-20} + \hat{\beta}_{t-20}(t - 4) + \hat{y}_{t-4|t-20}^M$$

to get the machine forecast of 20-quarter forward annual earnings log growth as

$$\mathbb{E}_{t-20} [\log (\text{earnings}_t^M) - \log (\text{earnings}_{t-4}^M)] = \hat{\beta}_{t-20}4 + \hat{y}_{t|t-20}^M - \hat{y}_{t-4|t-20}^M.$$

An alternative is to use the machine inputs to directly forecast 20-quarter forward annual earnings log growth $\mathbb{E}_{t-20} [\log (\text{earnings}_t^M) - \log (\text{earnings}_{t-4}^M)]$.

¹⁰A detailed description of the series is given in the Data Appendix of the online supplementary file at www.sydneyludvigson.com/s/ucc_data_appendix.pdf

Inflation predictor variables For y_j equal to inflation, the forecasting model considers the following variables:

1. $\mathbb{F}_{jt-k}^{(i)}[y_{jt+h-k}]$, lagged values of the i th type's forecast, where $k = 1, 2$
2. $\mathbb{F}_{jt-1}^{(s \neq i)}[y_{jt+h-1}]$, lagged values of other type's forecasts, $s \neq i$
3. $var_N \left(\mathbb{F}_{t-1}^{(\cdot)}[y_{jt+h-1}] \right)$, where $var_N(\cdot)$ denotes the cross-sectional variance of lagged survey forecasts
4. $skew_N \left(\mathbb{F}_{t-1}^{(\cdot)}[y_{jt+h-1}] \right)$, where $skew_N(\cdot)$ denotes the cross-sectional skewness of lagged survey forecasts
5. Trend inflation measured as $\bar{\pi}_{t-1} = \begin{cases} \rho \bar{\pi}_{t-2} + (1 - \rho) \pi_{t-1}, \rho = 0.95 & \text{if } t < 1991:\text{Q4} \\ \text{CPI10}_{t-1} & \text{if } t \geq 1991:\text{Q4} \end{cases}$, where CPI10 is the median SPF forecast of annualized average inflation over the current and next nine years. Trend inflation is intended to capture long-run trends. When long-run forecasts of inflation are not available, as is the case pre-1991:Q4, we use a moving average of past inflation.
6. \dot{GDP}_{t-1} = detrended gross domestic product, defined as the residual from a regression of GDP_{t-1} on a constant and the four most recent values of GDP as of date $t - 8$. See Hamilton (2018).
7. \dot{EMP}_{t-1} = detrended employment, defined as the residual from a regression of EMP_{t-1} on a constant and the four most recent values of EMP as of date $t - 8$. See Hamilton (2018).
8. $N_t^{(i)}[\pi_{t,t-h}]$ = Nowcast as of time t of the i th percentile of inflation over the period $t - h$ to t .

Lags of the dependent variable:

1. $y_{t-1,t-h-1}$ one quarter lagged inflation.

The factors in $\hat{\mathbf{G}}'_{jt}$ include factors formed from three large datasets separately:

1. $\mathbf{G}_{M,t-k}$, for $k = 0, 1$ are factors formed from a real-time macro dataset \mathcal{D}^M with 92 real-time macro series; includes both monthly and quarterly series, with monthly series converted to quarterly according to the method described in the data appendix.
2. $\mathbf{G}_{F,t-k}$, for $k = 0, 1$ are factors formed from a financial data set \mathcal{D}^F with 147 monthly financial series.

3. $\mathbf{G}_{D,t-k}^Q$, for $k = 0$ are quarterly factors formed from a daily financial dataset \mathcal{D}^D of 87 daily financial indicators. The raw daily series are first converted to daily factors $\mathbf{G}_{D,t}(\mathbf{w})$ and the daily factors are aggregated up to quarterly observations $\mathbf{G}_{D,t}^Q(\mathbf{w})$ using a weighted average of daily factors, with the weights \mathbf{w} dependent on two free parameters that are chosen to minimize the sum of squared residuals in a regression of $y_{t,t+h}$ on $\mathbf{G}_{D,t}(\mathbf{w})$.
4. *Macro data surprises* from the money market survey. The macro news include, GDP growth (Q/Q percentage change), core CPI (Month/Month change), unemployment rate (percentage point), and nonfarm payroll (month/month change). We include first release, second release, and final release for GDP growth. This constitutes six macro data surprises per quarter.
5. *FOMC surprises* are defined as the changes in the current-month, 1, 2, 3, 4, 6, 12, and 24 month-ahead federal funds futures (FFF) contract rate and the changes in the 1, 2, 3, 4, and 8 quarter-ahead Eurodollar (ED) futures contracts, from 10 minutes before to 20 minutes after each FOMC announcement. When benchmarking against a survey, we use the last FOMC meeting before the survey deadline to compute surprises. For surveys that do not have a clear deadline, we compute surprises using from the last FOMC in the first month of the quarter. When benchmarking against moving average, we use the last FOMC meeting before the end of the first month in each quarter to compute surprises. This leaves 10 FOMC surprise variables per quarter.
6. *Stock market jumps* are accumulated 30-minute window negative and positive jumps in the S&P 500 around news events over the previous quarter.

The 92 macro series in \mathcal{D}^M are selected to represent broad categories of macroeconomic time series. The majority of these are real activity measures: real output and income, employment and hours, consumer spending, housing starts, orders and unfilled orders, compensation and labor costs, and capacity utilization measures. The dataset also includes commodity and price indexes and a handful of bond and stock market indexes, and foreign exchange measures. The financial dataset \mathcal{D}^f is an updated monthly version of the of 147 variables comprised solely of financial market time series used in Ludvigson and Ng (2007). These data include valuation ratios such as the dividend-price ratio and earnings-price ratio, growth rates of aggregate dividends and prices, default and term spreads, yields on corporate bonds of different ratings grades, yields on Treasuries and yield spreads, and a broad cross-section of industry, size, book-market, and momentum portfolio equity returns.¹¹ The 87 daily financial indicators in \mathcal{D}^D include daily time series on commodities spot prices and futures prices, aggregate stock market indexes, volatility indexes, credit spreads and yield spreads, and exchange rates.

¹¹A detailed description of the series is given in the Data Appendix of the online supplementary file at www.sydneyludvigson.com/s/ucc_data_appendix.pdf

Employment growth predictor variables and specifications For earning growth forecasts, I first detrend the (log) employment level in real time by, starting with an initial sample, recursively running the following regression at each point in time t

$$\log(\text{employment}_t) = \alpha_t + \beta_t t + y_t \quad (\text{A.151})$$

For y_t equal to the detrended (log) employment level, I construct a forecasted value for y_t , denoted $\hat{y}_{t|t-h}$, based on information known up to time t using the following variables:

1. $\mathbf{G}_{M,t-k}$, for $k = 0, 1$ are factors formed from a real-time macro dataset \mathcal{D}^M with 92 real-time macro series; includes both monthly and quarterly series, with monthly series converted to quarterly according to the method described in the data appendix.
2. $\mathbf{G}_{F,t-k}$, for $k = 0, 1$ are factors formed from a financial data set \mathcal{D}^F with 147 monthly financial series.
3. $\mathbf{G}_{D,t-k}^Q$, for $k = 0$ are quarterly factors formed from a daily financial dataset \mathcal{D}^D of 87 daily financial indicators. The raw daily series are first converted to daily factors $\mathbf{G}_{D,t}(\mathbf{w})$ and the daily factors are aggregated up to quarterly observations $\mathbf{G}_{D,t}^Q(\mathbf{w})$ using a weighted average of daily factors, with the weights \mathbf{w} dependent on two free parameters that are chosen to minimize the sum of squared residuals in a regression of y_t on $\mathbf{G}_{D,t}(\mathbf{w})$.
4. *LDA factors* $F_{k,t-j}$, for topic $k = 1, 2, \dots, 50$ and $j = 0, 1$. The value of the topic k at time t is the average weights of topic k of all articles published at t .
5. *Macro data surprises* from the money market survey. The macro news include, GDP growth (Q/Q percentage change), core CPI (Month/Month change), unemployment rate (percentage point), and nonfarm payroll (month/month change). I include first release, second release, and final release for GDP growth. This constitutes six macro data surprises per quarter.
6. *FOMC surprises* are defined as the changes in the current-month, 1, 2, 3, 4, 6, 12, and 24 month-ahead federal funds futures (FFF) contract rate and the changes in the 1, 2, 3, 4, and 8 quarter-ahead Eurodollar (ED) futures contracts, from 10 minutes before to 20 minutes after each FOMC announcement. When benchmarking against a survey, I use the last FOMC meeting before the survey deadline to compute surprises. For surveys that do not have a clear deadline, I compute surprises using from the last FOMC in the first month of the quarter. When benchmarking against moving average, I use the last FOMC meeting before the end of the first month in each quarter to compute surprises. This leaves 10 FOMC surprise variables per quarter.
7. *Stock market jumps* are accumulated 30-minute window negative and positive jumps in the S&P 500 around news events over the previous quarter.

8. y_{t-k} for $k = 1, 2$ are lags of the dependent variable.
9. $\bar{\mu}_{t-k}$ for $k = 1, 2$ is the historical mean of the dependent variable calculated up to time t . The initial period is 1959Q1.
10. $\mathbb{F}_{t-k}[y_{t+h-k}]$ for $k = 1, 2$ are lags of the survey forecasts of the dependent variable.
11. For LSTM forecasts only: EN forecast of the variable to be predicted. Missing EN forecasts during earlier periods are imputed with the historical mean of the dependent variable.

The 92 macro series in \mathcal{D}^M are selected to represent broad categories of macroeconomic time series. The majority of these are real activity measures: real output and income, employment and hours, consumer spending, housing starts, orders and unfilled orders, compensation and labor costs, and capacity utilization measures. The dataset also includes commodity and price indexes and a handful of bond and stock market indexes, and foreign exchange measures. The financial dataset \mathcal{D}^f is an updated monthly version of the of 147 variables comprised solely of financial market time series used in Ludvigson and Ng (2007). These data include valuation ratios such as the dividend-price ratio and earnings-price ratio, growth rates of aggregate dividends and prices, default and term spreads, yields on corporate bonds of different ratings grades, yields on Treasuries and yield spreads, and a broad cross-section of industry, size, book-market, and momentum portfolio equity returns.¹² The 87 daily financial indicators in \mathcal{D}^D include daily time series on commodities spot prices and futures prices, aggregate stock market indexes, volatility indexes, credit spreads and yield spreads, and exchange rates.

After I obtain the machine forecast for the detrended level of employment, y , I obtain the h -horizon machine employment growth forecast (from $t-h$ to t denoted $\mathbb{E}_{t-h} [\Delta \log (\text{employment}_t^M)]$) by constructing

$$\mathbb{E}_{t-h} [\Delta \log (\text{employment}_t^M)] = \hat{\alpha}_{t-h} + \hat{\beta}_{t-h}t + \hat{y}_{t|t-h}^M - \log (\text{employment}_{t-h}) \quad (\text{A.152})$$

where $\log (\text{employment}_{t-h})$ is the realized log employment level at time $t-h$, and $\hat{y}_{t|t-h}^M$ is the machine forecast of the detrended log employment based on information up to time $t-h$. Following De La O and Myers (2021) and Bordalo et al. (2019), I assume that the survey respondents and the machine could observe $\log (\text{employment}_{t-h})$, the log employment level for the quarter on which the forecast is made.

¹²A detailed description of the series is given in the Data Appendix of the online supplementary file at www.sydneyludvigson.com/s/ucc_data_appendix.pdf