

Job Search Technology, Selection, and Hiring *

Sadhika Bagga
University of Rochester

Lukas F. Mann
Arizona State University

Version: October 16, 2026

Abstract

Online job search technology has increased the volume of job applications since the 1990s. We study the labor market effects of this shift. In time series and cross-sectional data, the rise of online job search is associated with signs of more selective firm screening: vacancies take longer to fill but are filled with more positively selected workers; for those not selected, unemployment becomes a more persistent state. We develop a labor market model with worker heterogeneity and screening in which more applications incentivize more selective screening. The decrease in application costs can quantitatively account for the recent fall in matching rates and lengthening of unemployment spells. We use the model to evaluate the impact of artificial intelligence on matching and find that it has accelerated trends set in motion by online job search.

*Bagga: University of Rochester, sadhika.bagga@rochester.edu, www.sadhikabagga.github.io;
Mann: Arizona State University, lfmann@asu.edu, www.lukasmann.com. We thank Brandon Miskanic for excellent research assistance. We thank seminar participants at Arizona State University, Cornell University, Duke University, the Federal Reserve Bank of San Francisco, the Federal Reserve Bank of St. Louis, Stony Brook University, the University of Rochester, and conference participants at the 3rd HKU Macroeconomics Workshop, the 5th DC Search and Matching Conference, the Sciences Po Search and Matching Workshop, the Women in Macro Conference at the University of Chicago, the BSE Summer Forum, the SED Meetings, and the NBER Summer Institute. A conference draft of this paper was circulated with the title "A Theory of Congestion and Rising Unemployment Durations."

Over the past three decades, the technology through which workers and firms meet in the labor market has advanced substantially. Before the mid-1990s, job search relied on newspaper classifieds and mailed or faxed resumes. With the rise of the internet, classifieds migrated to online job boards, enabling job seekers to apply to positions instantly. More recently, the emergence of artificial intelligence tools that generate application materials at scale has further reduced the cost and effort of applying. Accompanying these technological shifts, the volume of applications submitted per job seeker has risen sharply. The literature on frictional labor markets has long recognized that such improvements in search technology can improve the quality of matches or increase the rate at which they are formed (Stigler, 1961, 1962; Martellini and Menzio, 2020).

In this paper, we study a novel channel through which a higher application volume affects the labor market: it incentivizes firms to become more selective in whom they hire. Through this channel, advances in job search technology can slow down the formation of matches and make job finding prospects more unequal. We show that the US labor market has evolved in line with this view in the last three decades: as the volume of applications rose, firms slowed down hiring, devoted more resources to screening, and shifted hiring toward a more positively selected set of workers. For those repeatedly screened out, unemployment spells lengthened substantially.

Our first contribution is to document these novel facts and show that they are more pronounced in sectors and establishments that rely more on online job search. Our second contribution is to develop a model of the frictional labor market with heterogeneous workers, applications and screening to formalize the mechanism. We use the model to study the aggregate and distributional effects of online job search and find that changes in job search technology can explain the concentrated rise in unemployment duration, fall in matching rates, and change in the composition of hires we observe since the late 1990s. Finally, we show that our model offers a useful framework to understand the effects of artificial intelligence on labor market matching.

We begin by documenting four developments in the US labor market that are consistent with the mechanism we posit. First, we use data from the Employment Opportunity Pilot Project (EOPP) and Survey of Consumer Expectations (SCE) to show that the increasing use of the internet in job search coincided with a strong rightward shift in the distribution of applications sent by unemployed job seekers. Second, we use data from the Job Openings and Labor Turnover Survey (JOLTS) and vacancy postings from Lightcast to document that the rate at which jobs were filled declined during this time, while firms increased their demand for human resources personnel tasked with screening job appli-

cants. Third, we build on the methodology of [Mueller \(2017\)](#) to show that the average quality of hires in the Current Population Survey (CPS) exhibited an upward trend over this horizon: hires increasingly came from workers with higher pre-separation wages and relatively lower unemployment durations. Finally, the average length of unemployment spells nearly doubled over this period. Using data from the National Longitudinal Survey of Youth (NLSY), we show that this increase was concentrated among a subset of workers characterized by longer average spells and lower Armed Forces Qualification Test (AFQT) scores.

Next, we examine how these patterns are associated with online job search in cross-sectoral and cross-establishment data. We construct a sector-level measure of online job search exposure based on the sector's share of vacancies posted online compared to its overall vacancy share in JOLTS. We show that sectors with a higher prevalence of online job search experienced larger declines in the job filling rate as well as larger increases in screening personnel and in unemployment duration. Lastly, we document using micro-level survey data from Germany that establishments that utilize online recruitment methods receive more applications, screen applicants more intensely, take longer to fill their vacancies, and are less likely to hire workers with longer unemployment durations.

We argue that these findings jointly establish evidence of a new mechanism through which the rise of online job search affected labor market matching: as firms faced an increasing number of applications, they became less willing to hire applicants without obtaining more information about their unobserved productivity. This led firms to screen more intensely and to slow down hiring. As a result, the composition of hires shifted toward more productive workers, while less productive workers became increasingly locked out of employment and, over time, started experiencing substantially longer unemployment spells.

To formalize this mechanism, we develop a general equilibrium model of the labor market in which firms make optimal hiring decisions by trading off expected vacancy duration against expected worker productivity. Our model is a parsimonious extension of the canonical Diamond-Mortensen-Pissarides (DMP) framework, to which we add three ingredients. First, we introduce a notion of applications: after encountering a vacancy, workers choose whether to submit a costly application. Second, we introduce a notion of screening: workers differ permanently in productivity, which firms do not initially observe; before deciding whether to hire, a firm can pay to receive a signal about an applicant's type. Third, vacancy creation costs are sunk, which generates a trade-off between hiring the current applicant and waiting for a better one. Despite this added complexity,

the model remains tractable enough to be characterized graphically and features intuitive comparative statics.

In the model, a lower cost of applying leads workers to submit more applications, which, other things equal, raises their job finding rates. Since higher job finding rates improve the worker's outside option, their bargained wage increases. For the firm, this wage increase makes it less worthwhile to hire applicants of low expected productivity. Firms instead prefer to keep vacancies open and wait for applicants with higher expected productivity. This response more than offsets the first order effect of higher application rates and therefore leads to a net reduction in job finding rates. The decrease disproportionately affects less productive applicants who start experiencing increasingly persistent unemployment spells.

Taking our model to the data, we show that the decline in application costs brought about by online job search generates aggregate and distributional outcomes that are consistent with the data. In the aggregate, the model quantitatively matches a modest decline in the job finding rate alongside a substantial increase in the average duration of unemployment spells since the late 1990s. It also generates the observed decline in both the job filling rate and aggregate matching efficiency over the same time period. The model's distributional effects are likewise consistent with the data: the composition of hires shifts toward workers with higher pre-separation wages and relatively shorter unemployment durations, consistent with our evidence from the CPS. Some workers experience longer unemployment spells on average, while others are largely unaffected, in line with our evidence from the NLSY. Lastly, the growing concentration of workers with low job finding rates among the long-term unemployed generates stronger observed duration dependence that we document in the CPS.

In the final part of the paper, we use our framework to study the equilibrium effects of more recent advances in search technology that, besides the cost of applying, affected the cost and quality of screening.¹ Specifically, we study how the deployment of artificial intelligence (AI) tools in job search affected labor market matching. We account for three separate channels: (i) job seekers use AI tools to generate application materials at scale, which lowers the cost of applying;² (ii) firms deploy AI-based applicant tracking systems

¹The reason for the focus on applications in the context of online job search is measurement. The number of applications submitted by job seekers is observed both before and after the proliferation of online job search, whereas comparable evidence on how the cost and the quality of screening evolved over the same time period is not available.

²See Wall Street Journal, [“You’re Fighting AI With AI: Bots Are Breaking the Hiring Process,”](#) accessed September 23, 2026.

which lowers the cost of obtaining signals about applicants;³ and (iii) AI-generated applications make it harder to distinguish between qualified and unqualified candidates, which lowers the informativeness of the signals firms obtain (Cowgill et al., 2024; Galdin and Silbert, 2025). We show that, in our model, the three channels affect all aggregates through a single sufficient statistic: cheaper applications and cheaper signals make firms more selective, while less informative signals have the opposite effect. We discipline each channel using available evidence and find that the first two dominate. The net effect is therefore that the deployment of AI tools accelerates the developments set in motion by online job search: it accounts for about three-fifths of the decline in the job finding rate and four-fifths of the decline in the job filling rate observed since 2019.

Related Literature. Our paper contributes to three growing strands of literature: the literature on the effects of diminishing search frictions on labor market outcomes, the literature on the causes of long-term unemployment, and an emerging literature on the effects of artificial intelligence on labor market matching.

The idea that costly search is a key determinant of labor market frictions dates back to Stigler (1961, 1962), and the question of how improvements in search technology affect these frictions has received renewed attention in the context of the internet revolution. Bhuller et al. (2023) provide quasi-experimental evidence on how the staggered rollout of broadband infrastructure in Norway affected labor market outcomes. On the theoretical side, Martellini and Menzio (2020) study the long-run consequences of declining search frictions in a model where firms screen matches based on idiosyncratic match quality. In their framework, lower frictions improve match selection but leave unemployment transition rates unchanged. Birinci et al. (2025) reach a similar conclusion: when application rates increase but firms screen on match quality, the main effect is a decline in the separation rate while the job finding rate remains roughly unchanged. In both papers, declining search frictions thus lead to higher match quality, while matching rates stay roughly level. In contrast, we highlight several distinct implications that arise when firms screen on permanent worker characteristics instead of idiosyncratic match quality: a positive trend in the composition of hires, a decrease in aggregate matching efficiency, and an uneven lengthening of unemployment spells. Auster et al. (2025) highlight a related mechanism in a directed search framework with adverse selection, showing that declining search frictions can worsen matching outcomes. We obtain a similar result in

³See Business Insider, “This Startup Is Replacing Job Recruiters With AI Agents,” accessed September 23, 2026.

a random search setting that is designed to closely resemble the canonical DMP model, which allows us to link the mechanism to empirical labor market trends.^{4,5}

Our paper also contributes to the literature on the causes of increasing unemployment duration. [Mukoyama and Şahin \(2009\)](#) emphasize the role of rising wage dispersion in driving the trend, arguing that it raised reservation wages and lengthened unemployment spells through increased selectivity.⁶ The literature has also studied the role of duration dependence in shaping the duration of unemployment spells over the business cycle; that is, the idea that unemployment exit hazards decline as unemployment progresses ([Kroft et al., 2016](#)). More recent work has emphasized that duration dependence may reflect dynamic selection ([Jarosch and Pilossoph, 2019](#); [Mueller et al., 2021](#); [Ahn, 2023](#)), a view our theory adopts. We offer a new perspective on the relationship between duration dependence and unemployment duration. Increasing firm selectivity widens the *dispersion* of job finding rates across workers, which steepens observed duration dependence and rationalizes a large overall increase in the average duration of unemployment even when job finding rates decline only moderately.

Finally, we contribute to a rapidly growing literature on the effects of artificial intelligence on matching markets. [Cowgill et al. \(2024\)](#) and [Galdin and Silbert \(2025\)](#) show that access to generative AI lowers the information content of job applications. Related work studies additional dimensions of how AI reshapes hiring (see, e.g., [Gans, 2024](#); [Cui et al., 2025](#); [Dhillon et al., 2025](#); [Wiles and Horton, 2025](#)). We contribute to this literature by conducting a general equilibrium analysis that jointly accounts for three channels emphasized in this empirical work: AI makes it easier for workers to contact employers, it makes it easier for employers to screen; and it makes it harder for employers to distinguish between applicants.

The remainder of the paper proceeds as follows. Section 1 documents the evolution of the US labor market over the last three decades and links the observed trends to online job search in cross-sectional and establishment-level data. Section 2 formulates the model and uses it to study the effects of a decline in application costs. Section 3 shows the quantitative fit of these effects at the micro and macro level. Section 4 uses the model to

⁴Our model integrates applications and a screening process into a standard DMP framework, joining a number of papers that have recently incorporated these features into search models (see, e.g., [Wolthoff, 2018](#); [Engbom, 2021](#); [Bradley, 2025](#); [Cai et al., 2025](#)).

⁵[Pries and Rogerson \(2022\)](#) is another notable example of a screening model with idiosyncratic match productivity although their focus is on explaining the secular decline in worker turnover and short-duration employment spells.

⁶We study data on reservation wages in Appendix A.4. We find no long-term increase in reservation wages but show that, consistent with our theory, the dispersion of reservation wages has widened.




Figure 1: Proliferation of online job search and surge in job applications

Source: Current Population Survey Computer and Internet Use Supplement, Employment Opportunity Pilot Project (1979-80), Survey of Consumer Expectations (2013-22). *Notes:* The left panel plots the fraction of unemployed job seekers who reported using the internet to search for jobs. The right panel plots the distribution of applications submitted in the prior four weeks by unemployed job seekers between 1979-2022.

study the impact of artificial intelligence on labor market matching. Section 5 concludes.

1 Evolution of the US labor market

1.1 Online job search and job applications

Over the past three decades, the internet has transformed the technology through which workers search for jobs and firms search for workers. The left panel of Figure 1 illustrates this development by plotting the share of unemployed job seekers who reported using the internet for job search in the CPS Computer and Internet Use Supplement.⁷ This share stood at 40% in 1998 and rose steeply over the following decade, coinciding with the emergence and rapid growth of major online job boards such as Monster.com, LinkedIn and Indeed. By the mid-2010s, internet-based job search had become ubiquitous, used by over 70% of the unemployed and plateauing at that level through the end of the sample.

The shift to online search was accompanied by a surge in the number of job applications. The right panel of Figure 1 plots the distribution of applications submitted by

⁷The question “Do you use the Internet to search for a job?” was administered at irregular intervals in the CPS surveys and the dashed segments of the left panel connect non-consecutive question years.




Figure 2: Jobs take longer to fill and firms increase demand for screening

Source: Current Population Survey, Job Openings and Labor Turnover Survey, [Petrosky-Nadeau and Zhang \(2021\)](#) and Lightcast. Notes: The left panel plots the job filling rate which is the ratio of monthly hires and vacancies, averaged quarterly. The right panel plots online job postings for Human Resources Specialists (SOC 13-1071) as a share of postings in Business and Financial Operations occupations (SOC 13).

unemployed job seekers over the preceding four weeks, combining data from the Employment Opportunity Pilot Project (EOPP) for 1979–80 with the Survey of Consumer Expectations (SCE) for 2013–22.⁸ We report details of both data sets and the comparability of job application measures in Appendix A.1. In 1979–80, applications were concentrated at the bottom of the distribution: two-thirds of unemployed applicants submitted between one and three applications, while only 2% submitted twenty or more. By 2021–22, the distribution had shifted to the right and these shares changed to 28% and 21%, respectively.

The increasing use of the internet for job search and the concurrent rise in application counts suggest that the emergence of online job search lowered the cost of submitting job applications and therefore increased the rate at which workers applied to jobs. In what follows, our goal is to understand how the labor market responded to this surge in applications. We first turn our attention to changes in firm-side labor market variables over this period.

1.2 Job filling rate and the demand for screening

The surge in applications after the 1980s did not translate into faster hiring; instead, vacancy filling rates declined. The left panel of Figure 2 plots the job filling rate—the ratio of

⁸To our knowledge, no comparable data on application counts is available for the US between 1980 and 2013.

monthly hires to the stock of vacancies—between 1975 and 2025.⁹ The series was highly cyclical but displayed no strong trend from 1975 through the early 2000s, before shifting sharply downward after the Great Recession. Between the late 1990s and late 2010s, two periods with unemployment near 4%, the job filling rate fell from 1.75 to 1.40, a decline of 20%.¹⁰ Inverting the job filling rate expresses this decline in terms of time to fill. Whereas in the 1970s and 1980s a vacancy generated a hire within two to three weeks, by the 2020s it took more than a month.¹¹

Why did the job filling rate decline over a period when firms attracted a larger number of applications? One explanation is that firms used the larger applicant pool to become more selective and screen applicants more intensely, instead of hiring workers more rapidly. To test whether the data is consistent with this hypothesis, it is useful to examine how the time and effort that firms spent on screening applicants has evolved over time. In the US, direct measures of such screening intensity are not available. We therefore proxy for it by tracking the six-digit occupation “Human Resources (HR) Specialists” whose task, as defined by the Bureau of Labor Statistics, is to “recruit, screen, interview, or place individuals within an organization”. The occupation pertains exclusively to the recruitment and evaluation of job candidates, and excludes adjacent HR functions such as payroll processing, personnel training, etc.

The right panel of Figure 2 plots job ads for HR specialists as a share of job ads for business and financial operations occupations, the two-digit occupation family to which the six-digit HR Specialists belong.¹² The series exhibits an upward trend, rising from 12.5% in 2007 to nearly 17% by 2019, indicating an increasing firm demand for labor used to screen job applicants. We also document in Appendix A.2 that the employment share of HR specialists displays a similar pattern over a longer horizon: it rose from about 6% in 2012 to nearly 9% in 2024.

In sum, the evidence presented above shows that, compared to the pre-internet era, job seekers now submit more applications on average, yet vacancies take longer to fill and firms demand more personnel whose task is to screen job applicants. These patterns

⁹We measure vacancies from the JOLTS and hires from the CPS and describe details of construction in Appendix A.2

¹⁰In Appendix A.7, we plot the job filling rate against labor market tightness and unemployment rate to demonstrate how it varied over time conditional on the state of the business cycle. We note that the job filling rate decreased conditional on the state of the business cycle starting from the Great Recession.

¹¹This is consistent with Faberman and Menzio (2018), who use data from the Employment Opportunities Pilot Project and document that it took on average 16.2 days to fill a vacancy in 1982. We are not aware of contemporaneous source of vacancy duration for the US.

¹²The data on job ads come from Lightcast.

are consistent with the idea that firms respond to the larger applicant pool by screening more intensely and becoming more selective in who they hire. This interpretation yields a natural testable implication: if screening made hiring more selective, the composition of hires should shift toward more positively selected workers, while workers who are less sought-after by firms should remain persistently unemployed. We next examine whether evidence is consistent with this implication.

1.3 Shifting selection of hires

We begin by measuring the composition of hires over time. To do so, we use two complementary rank measures constructed from the CPS. The first builds on [Mueller \(2017\)](#) and asks where newly hired workers from non-employment ranked in the wage distribution *before* they became nonemployed. To construct this measure, we utilize hourly wages of employed workers which are recorded in the fourth month of a worker's tenure with the CPS. We then assign ranks to each worker based on the position of their wage record relative to the wage distribution of that year. Next, we track the worker's labor force status during their final four interview months with the CPS, roughly one year after their wage is observed. Based on these observations of labor market status, we compute the mean rank of hires from non-employment, that is, workers who are employed in their final interview but who are nonemployed in at least one of the three preceding months. We call this the pre-displacement wage rank of hires.

Our second measure of the quality of hires asks where newly hired workers ranked in the unemployment duration distribution at the time they were hired. For each unemployed worker we compute the rank of their elapsed unemployment duration relative to all other unemployed workers in that year. The duration rank is the average of this rank among unemployed workers who transition into employment in the following month. To ensure that both measures capture compositional shifts in unobservables rather than demographics, we residualize individual wage and duration measures with respect to occupation, education, and age, before computing ranks. In [Appendix A.3](#), we provide more details of the construction of both measures.

[Figure 3](#) shows that new hires have become increasingly positively selected over time, based both on duration and prior wages. The pre-displacement wage rank, plotted in the left panel, rose modestly through the early 1990s, and rose more strongly thereafter, standing at about five percentile points above its initial level at the end of the sample. New hires thus increasingly shifted toward workers who had earned relatively high wages in

(a) Pre-displacement wage rank of new hires (b) Unemployment duration rank of new hires




Figure 3: New hires positively selected on wage and unemployment duration ranks

Source: Current Population Survey. *Notes:* The left panel plots the mean pre-displacement wage rank of hires from non-employment. The right panel plots the mean elapsed unemployment duration rank of hires from unemployment. Ranks are computed within each year after residualizing log wages and log duration with respect to age, education, and occupation. Both series are centered three-year moving averages, normalized to zero at the beginning of their sample.

their previous jobs. The right panel shows that the unemployment duration rank of new hires decreased between the 1990s and the mid-2020s, ending about four percentile points below its initial level. This indicates that hires toward the end of the sample were increasingly drawn from workers with relatively shorter unemployment spells. Both patterns point to an increasingly positive selection of hires over time: the composition of hires has shifted toward workers with stronger pre-match labor market outcomes.

1.4 Evolution of unemployment duration

How did this shift in the composition of hires affect the job finding prospects of the unemployed? Our evidence on hires suggests that, over time, firms increasingly selected workers with stronger labor market attributes into employment. A plausible worker-side effect of this trend is that workers with weaker attributes may find it increasingly difficult to find work and experience longer periods of joblessness. We next turn to two pieces of evidence that substantiate this conjecture: the average duration of unemployment spells increased and the rise was concentrated among a subset of workers.

Figure 4 plots the average length of ongoing unemployment spells between 1975 and 2025. While the average duration of ongoing spells was 2.5 months during the expansions of the late 1970s and 3.2 months in the late 1990s, by the late 2010s, it had reached 5.3




Figure 4: Mean unemployment duration

Source: Bureau of Labor Statistics based on Current Population Survey. *Notes:* The figure plots the centered twelve-month moving average of the mean elapsed duration of ongoing unemployment spells.

months and has remained near that level since. Underlying this rise is a growing prevalence of long-term unemployment: the share of job seekers unemployed for six months or longer grew from about 8% in the late 1970s to about 28% at the end of the sample (Appendix Figure A.4).¹³


If more selective hiring based on permanent worker characteristics was driving the rise in average duration, the burden should fall unevenly on workers: longer unemployment spells should be concentrated among those workers who are increasingly screened out, while those who are preferred by firms should be less affected. We test this hypothesis using panel data from the National Longitudinal Surveys of Youth (NLSY). The NLSY follows two cohorts of individuals entering the labor force in 1979 and 1997, respectively, from their labor market entry to the present day. Each worker's labor force status is recorded at a weekly frequency.

Within each cohort, we classify workers based along two dimensions. First, their lifetime average duration of completed unemployment spells; we compare workers above the median with those below the median. Second, by their overall score in the Armed Forces Qualification Test (AFQT); we compare workers in the first quintile with those in the fifth.¹⁴ The two panels of Figure 5 plot the average duration of unemployment spells

¹³In Appendix A.4, we discuss competing explanations for the rise in average duration, such as the shifts in the composition of the labor force over time, changes in unemployment insurance, and changing reservation wages among workers. We show that they cannot by themselves account for the sustained increase in unemployment duration.

¹⁴Since AFQT scores are a noisier predictor of labor market desirability than lifetime average unemployment duration, we choose to compare the top and bottom quintile instead of separating along the median.

(a) By lifetime average unemployment duration



(b) By AFQT scores




Figure 5: Mean unemployment duration by two measures of worker classification

Source: National Longitudinal Surveys of Youth 1979 and 1997. Notes: Each panel plots the mean duration of completed unemployment spells among workers averaged over two-year bins. Red lines refer to the NLSY79 cohort and blue lines to the NLSY97 cohort. In the left panel, the dashed lines (“above median”) refer to workers whose lifetime average unemployment duration is above the cohort median, solid lines (“below median”) refer to those at or below it. In the right panel, the dashed lines (“low AFQT”) refer to workers in the bottom AFQT quintile and solid lines (“high AFQT”) to workers in the top AFQT quintile.

along each classification over time.

The rise in unemployment duration is driven almost entirely by workers with longer lifetime average unemployment durations and lower AFQT scores: for workers with above median lifetime average unemployment duration, the average rises from approximately four to eleven months for the 1979 cohort, and three to twelve months for the 1997 cohort. Average duration among workers below the median remained close to two months for both cohorts. The rise in aggregate unemployment duration thus reflects a growing dispersion of job finding prospects across workers.¹⁵ Likewise, we observe a divergence across groups over time when we classify workers by their AFQT scores. Duration rises more strongly among workers in the bottom AFQT quintile, although it also rises among those in the top quintile.

In sum, Figures 3 and 5 describe two sides of the same compositional shift. Hires increasingly come from workers with stronger pre-match labor market outcomes, while

In Appendix A.5, we describe additional details of construction.

¹⁵This observation is consistent with recent work that fits duration models to CPS data (Ahn and Hamilton, 2020; Ahn, 2023) and finds that increasing heterogeneity in job finding rates has contributed to a rise in unemployment duration. Furthermore, in Appendix A.7, we argue that the increase in unemployment duration larger than the decrease in the job finding rate is an additional piece of evidence that the dispersion in job finding rates increased over time.




Figure 6: Aggregate trends pronounced in sectors more exposed to online job search

Source: Job Openings and Labor Turnover Survey, Occupational Employment and Wage Statistics, Current Population Survey and Lightcast. *Notes:* Each circle represents a two-digit NAICS sector. The x-axis captures online exposure which is the log deviation of a sector's share of Lightcast postings to its share of JOLTS vacancies in 2019. The y-axis reports changes from 2000 to 2023: log changes in job filling rates and unemployment duration, and level changes in the HR-specialist employment share. Circle sizes are proportional to initial employment, and gray lines are employment-weighted linear fits.

longer unemployment spells become concentrated among workers with persistently weak job-finding prospects. This joint pattern is consistent with increasingly selective hiring over a heterogeneous set of workers. Our final piece of evidence, presented next, shows that these implications are empirically more prevalent in sectors and establishments that are more exposed to online job search.

1.5 Cross-sectional analysis

The evidence presented so far documents aggregate trends that unfolded alongside a shift to online job search. In this section, we consider whether the two are systematically related in the cross-section. We use two sources of cross-sectional variation: sectors that differ in their exposure to online job search, and micro-data on German establishments that differ in the search channels through which they hire.

1.5.1 Sectoral heterogeneity in online job search

We construct a measure of exposure to online job search for two-digit NAICS sector s as:

$$\text{Online exposure}_s = \log \left(\frac{V_s^{\text{online}}}{\sum_j V_j^{\text{online}}} \right) - \log \left(\frac{V_s}{\sum_j V_j} \right), \quad (1)$$

where V_s^{online} denotes online job postings from Lightcast, and V_s denotes total online and offline postings in JOLTS, both measured in 2019.¹⁶ A value of zero means that a sector accounts for the same share of online postings as of all vacancies, while a positive value means that it is disproportionately represented online. We find that exposure varies substantially across sectors. Educational Services accounts for about 3.4 times as large a share of online postings as of all vacancies, whereas the corresponding ratio for Wholesale Trade is less than one-third.

Figure 6 relates this measure to the long-run changes in three sectoral outcomes between 2000 and 2023. Each circle represents a sector, its size is proportional to the sector's initial employment, and the gray line is the employment-weighted linear fit. The left panel shows the change in the job filling rate plotted against online exposure. Sectors more exposed to online search experienced larger declines in their job filling rates. The middle panel plots the change in the employment share of HR specialists within business and financial operations occupations, which is our proxy for resources devoted to recruiting and screening. The relationship is positive: an increase in exposure is associated with an increase in the HR-specialist share. Lastly, the right panel shows a positive relationship between online exposure and the increase in average unemployment duration.¹⁷ Each of the aggregate developments documented in Sections 1.2–1.4 – slower vacancy filling, rising screening employment, and lengthening unemployment spells – is thus concentrated in the sectors where job search moved online.

Some of the data that would substantiate the mechanism we postulate, such as data on screening strategies employed by firms or the composition of hires, is unavailable at the

¹⁶The premise of the measure is that the nature of a sector's jobs determines whether its vacancies are posted online: an opening for a barista is often advertised at the door of the establishment, whereas an opening for a financial analyst is almost always posted online. In Appendix A.7, we show that job seekers in more exposed sectors submit more applications and saw larger increases in internet usage for job search. We measure exposure in 2019 to characterize which sectors online job search ultimately penetrated by the time the technology had matured.

¹⁷In Appendix A.6, we show that this relationship also holds when we compare US states rather than sectors. Moreover, we also show that matching efficiency and job finding rates also declined more in sectors that were exposed to online job search.

	Log applications (1)	Log days to fill (2)	Screening effort		Pr(hiring LTU) (5)
			Log screening hours (3)	Assessment methods (4)	
Searched online	0.815*** (0.021)	0.505*** (0.021)	0.668*** (0.027)	0.241*** (0.048)	−0.040* (0.018)
Observations	93,325	87,637	43,621	5,384	19,660
R^2	0.235	0.094	0.156	0.102	0.057

Table 1: Online search and recruitment outcomes in German establishments

Source: IAB Job Vacancy Survey, 2005-2021. *Notes:* Each column reports a separate regression of the dependent variable (shown in the column name) on an indicator for whether the establishment used at least one online search channel for its most recent hire. LTU denotes long-term unemployed. All specifications include year fixed effects and controls for the vacancy’s required skill level, broad occupations and industries, and the establishment’s log employment and log gross hiring rate. Outcomes are available in different survey waves (applications and days to fill 2005-21; screening hours 2014-15 and 2017-21, assessment methods 2018; and LTU hiring 2008-21) accounting for the differences in sample size. Standard errors, clustered by establishment, in parentheses. Establishments weights used. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

sector level. In what follows, we therefore supplement our evidence with establishment-level data.

1.5.2 Micro-evidence from German establishments

Since detailed establishment-level micro data on firm hiring and recruitment outcomes does not exist in the US, we turn to evidence from Germany. Specifically, we draw on the Job Vacancy Survey (JVS), administered by the German Institute for Employment Research (IAB). The JVS is a representative annual survey of a cross-section of German establishments that records, for the establishment’s most recent hire, the channels used to search for applicants, the number of applications received, the time taken to fill the position, the resources spent on screening candidates, and the characteristics of the worker hired. The survey spans 2005–2021 and covers roughly 14 thousand establishments per year.

We compare establishments that utilized online job search with those that only used offline methods. We classify a search as online if the establishment used at least one of the following methods: posting a job ad on its own website or an online job board, using the Federal Employment Agency’s websites, or using social media. Each column of Table 1 reports a different outcome of the recruitment process regressed on this indicator. The regression includes survey-year fixed effects and controls for the vacancy’s skill re-

quirement and occupation, the establishment's industry and size, and the establishment's recent hiring rate. The regressions describe correlations for each establishment's most recent hire. We report the construction of all the variables, their harmonization across survey waves, and unconditional correlations in Appendix A.6.

Table 1 shows that vacancies advertised online attract more than twice as many applications as observationally similar vacancies posted offline (column 1). Yet the larger applicant pool does not translate into the establishment reaching their hiring decision more quickly. Online vacancies take about 66% longer to reach their hiring decision from the beginning of search (column 2). What do employers do during this longer recruitment process? Establishments searching online devote more effort to evaluating candidates – working hours spent on recruiting and screening are nearly twice as high (column 3). This measure covers the entire recruitment process, from preparing the advertisement through evaluating candidates and finalizing the employment contract. Establishments that search online also deploy more methods to assess applicants beyond job interviews and application materials (column 4). These methods include conducting standardized tests, consulting former employers, doing background checks, etc. Finally, conditional on hiring from unemployment, establishments using online search methods are 4 percentage points less likely to hire a long-term unemployed worker, defined as someone whose unemployment spell has lasted for more than a year (column 5), consistent with the idea that more intensive screening translates into stronger selection.

Taking stock. We close the empirical section by summarizing the facts we documented. Over the last five decades: (i) job search moved online, and the number of job applications submitted rose sharply; (ii) vacancies took longer to fill, while firms demanded a growing share of personnel tasked with screening job candidates; (iii) new hires became increasingly positively selected based on pre-match labor market outcomes; and (iv) average unemployment duration rose and became concentrated among workers with persistently weak job-finding prospects. These aggregate developments were more pronounced in sectors with more online vacancies, and establishments hiring through online channels.

Taken together, these developments point toward the following mechanism: the increase in application rates incentivized firms to screen more intensely and become more selective in their hiring. This change had aggregate and distributional effects: in the aggregate, hiring rates per vacancy declined as firms adjusted their screening strategy. Distributionally, the shift in firm screening behavior led to a lengthening of unemployment spells that was concentrated among a subset of workers.

In the next section, we explore this mechanism through the lens of a search model that captures how firms optimally choose their screening strategy in a world with heterogeneous workers and information frictions. We use the model to study the relationship between declining application costs, endogenous screening, and aggregate labor market outcomes.

2 Model

The model is an extension of the canonical Diamond-Mortensen-Pissarides (DMP) model with three modifications. First, we introduce a notion of applications: upon meeting a vacancy, workers draw a random cost and decide whether to apply. Second, we introduce a notion of screening. There are two permanent, unobserved worker types who differ only in their productivity. Firms can pay a random cost to receive a signal about an applicant's type before making a hiring decision. Third, vacancy costs are sunk, which introduces a trade-off between filling a job immediately and waiting for a worker of higher expected productivity. Otherwise, the ingredients are those of a standard search-and-matching model.

We begin by outlining the environment (Section 2.1). We then describe the central decision problem of the model: a firm-side trade-off between filling vacancies quickly or filling them with better workers (Section 2.2). This allows us to characterize the equilibrium of the model graphically (Section 2.3) and to derive an intuition for the economic effects of a fall in application costs (Section 2.4) before taking the model to the data (Section 2.5).

2.1 Environment

Demographics. Time is continuous. Throughout the paper, we focus on stationary equilibria and thus suppress time subscripts. The economy is populated by a unit measure of infinitely-lived workers who can be either employed or unemployed. Workers permanently belong to one of two types $i \in \{L, H\}$ that differ in their productivity, with type- L workers comprising an exogenous share α_L^{pop} of the population. Employed workers produce output y_i ($y_H \geq y_L > 0$), receive a wage w_i , and do not engage in search. Unemployed workers receive a flow payoff $b_i = (1 - \rho)y_i$, where $\rho \in (0, 1]$, and search for a match in a frictional labor market. On the other side of the market are one-worker

vacant jobs that search for a worker until they are filled.¹⁸ Filled and vacant jobs are destroyed at rate δ . Firms and workers discount the future at rate r .

Job creation. Jobs are homogeneous and atomistic. There is an infinite supply of potential entrants who can create new vacant jobs (“vacancies”) at will. To create a new vacancy, an entering firm must pay a fixed cost, k . This means that the value of a vacant job V is equal to k in equilibrium.¹⁹

Meeting technology. Vacancies and workers encounter each other at random in a frictional labor market. Letting u denote the measure of unemployed workers and v the measure of vacancies, the total flow of meetings is $m = Zv^\omega u^{1-\omega}$, where Z is the aggregate meeting efficiency and ω is the elasticity of meetings with respect to vacancies. The aggregate meeting rates for vacancies and workers are thus

$$q = m/v, \quad \lambda = m/u = Z^{\frac{1}{1-\omega}} q^{\frac{\omega}{\omega-1}} \quad (2)$$

Match formation. Upon an encounter, but before production begins, the worker-firm pair engages in a match formation game that comprises three stages: an application stage, a screening stage, and a hiring stage. We discuss each in turn.

1. *Application stage.* Upon encountering the vacancy, the worker draws an application cost c from a distribution with cdf $F(c) = \gamma_a c^{\eta_a}$. If the worker pays the cost c , the firm receives their application and the game moves on to the screening stage; otherwise, both parties return to search. We refer to γ_a as the “ease of applying”: a higher value corresponds to a lower cost of the average application. We use p_i^{app} to denote the (endogenous) probability that a type- i worker submits an application, conditional on an encounter.
2. *Screening stage.* When a firm receives a worker’s application, it does not observe the worker’s type. The firm can learn about the worker’s type by operating a costly screening technology that reveals a signal (L or H). The signal equals the worker’s

¹⁸We will also refer to jobs as firms. In our economy, as in most search models, ownership across jobs is irrelevant for outcomes and each job operates independently.

¹⁹This assumption is a departure from the flow cost assumption that is more typical in search models. In such models, costless free entry drives the value of a vacancy to zero in equilibrium, and it is therefore never optimal for a firm to forego a match with positive surplus. In contrast, assuming a sunk vacancy creation cost creates a positive opportunity cost of forming a match, creating a trade-off between matching now and waiting for a better applicant.

true type with probability $s \geq 0.5$, making it weakly informative.²⁰ We call s the “signal quality”: a higher value of s corresponds to a more informative signal. Obtaining the signal entails a one-time cost κ , drawn when the firm receives an application from a distribution with cdf $G(\kappa) = \gamma_s \kappa^{\eta_s}$. We refer to γ_s as the “ease of screening”: a higher value corresponds to a lower cost of the average signal. Finally, we use p^{sig} to denote the (endogenous) probability that the firm decides to obtain a signal.²¹ The game moves on to the hiring stage regardless of whether the firm has chosen to obtain a signal.

3. *Hiring stage.* In the final stage of the game, the firm decides whether to hire the applicant. If the firm does not hire the applicant, both parties return to search. If the firm hires the applicant, it relinquishes the option to return to search, the match is formed, and the worker’s true type is revealed.²²

Production and wage setting. After the worker and the firm form a match, they produce flow output y_i of which the worker retains a wage w_i and the firm retains a flow profit of $y_i - w_i$. We assume that wages w_i are determined via Nash bargaining with worker bargaining weight β . Jobs continue to produce until they are destroyed; afterwards, the worker returns to unemployment.

While stylized, this setup parsimoniously captures the ideas discussed in the first part of the paper: the flow of applications is moderated by the cost of applying and firms can respond to changes in the flow of applications by adjusting their screening strategy, which changes the duration for which vacancies remain open and the composition of hires. The screening response of firms therefore has asymmetric effects across different

²⁰For tractability, we assume that the firm knows the aggregate distribution of applicants but cannot costlessly update their beliefs based on an individual worker’s unemployment history, departing from some recent work (Jarosch and Pilossoph, 2019; Bradley and Mann, 2024). One way to interpret this assumption is that the signal encodes the relevant information that the firm is able to extract from a worker’s background, but that extracting such information is costly.

²¹In Appendix B.1, we show that this technological assumption is equivalent to assuming that the firm controls the probability of receiving a signal p^{sig} subject to a deterministic convex screening cost.

²²The assumption that the firm gives up search when it hires a worker ensures that firms cannot use hiring as an alternative screening device by hiring a worker, revealing their type, and then firing L -types. An equivalent assumption is that re-posting or firing costs are sufficiently high to deter the firm from firing an L -type worker and re-posting the vacancy; for instance, a re-posting or firing cost of k ensures that this is the case. An alternative setup would be one in which firms learn about a worker’s type slowly over time and separate if their beliefs become sufficiently pessimistic, as in Baley et al. (2022). We argue that the distinction between both approaches is not central to the mechanism we emphasize.

types of workers and changes how quickly matches are formed in the aggregate labor market.

The application decision of the worker and the signal acquisition decision of the firm follow standard Bellman equations. However, since they do not add substantively to the intuition of the model, we relegate them to Appendix B.2. Instead, in what follows, we focus on developing an intuition for the hiring decision of the firm. This decision is at the heart of the model: it determines how firms respond to changes in the flow of applications and therefore governs how the composition of hires changes across the whole economy when applications become less costly. The following section outlines the key trade-offs involved in this decision.

2.2 The hiring decision of the firm

The firm makes a decision on whether to hire or reject a worker in three possible situations: after it receives an H -signal, after it receives an L -signal, or after it declines to obtain a signal. We will consider the three cases in turn.

Consider first the case in which the firm has received an H -signal. In this case, the firm chooses to hire the applicant. The intuition is simple: firms prefer filling vacancies more quickly and filling them with workers who are more likely to be H -types. Rejecting workers after receiving an H -signal lowers both the likelihood of the worker being an H -type and the filling rate and is thus a dominated strategy. We show a formal proof in Appendix B.3.

Consider next the case in which the firm has received an L -signal. In this case, the firm rejects the applicant. The key argument for this result is that a signal is valuable only because it affects the hiring decision of the firm. If a firm were to hire on an L -signal, the signal would be irrelevant to its hiring decision. In that case, it would not be optimal to purchase the signal in the first place. We again relegate the formal proof to Appendix B.3.

Lastly, consider the case in which the firm has received no signal. A firm in this state faces two options. The first option is to hire the worker. Using J_i to denote the value of a filled job with a worker of type i and using a_L to denote the share of type- L workers among applicants, this option yields an expected value of $a_L J_L + (1 - a_L) J_H$. The second option is to reject the worker, in which case the firm receives the value of a vacancy, which we denote by V . The firm thus solves the problem

$$\max_{R \in [0,1]} (1 - R) \cdot [a_L J_L + (1 - a_L) J_H] + R \cdot V. \quad (3)$$

where R denotes the endogenous probability that the firm rejects an applicant without a signal. This maximization problem encodes an important trade-off between hire quality and filling speed. By choosing a low value of R , the firm is more likely to hire a worker without a signal and therefore fills the vacancy more quickly but with a worker of lower expected productivity. Conversely, if the firm chooses a high value of R , it is more likely to hire only those workers for whom it receives an H -signal. As a result, the average hire will be of higher productivity but the vacancy will stay open for longer before it is filled. R therefore captures the degree to which the firm resolves this trade-off in favor of worker quality.²³

Since the problem in (3) is linear, it may feature a corner solution ($R = 0$ or $R = 1$). Even though, under some parameters, equilibria exist for which such a corner solution is optimal, we do not focus on this class of equilibria here. Instead we focus on *mixed-strategy equilibria*, for which R lies in the interior ($R \in (0, 1)$).²⁴ The reason for focusing on mixed-strategy equilibria is that the key trade-off described above operates through a smooth response of R to changes in application inflows, which is absent when the economy is in a pure-strategy equilibrium. We find that, in practice, there is a large parameter space that admits mixed-strategy equilibria with desirable properties. In particular, it is easy to find parameters that deliver mixed-strategy equilibria with moments in reasonable empirical ranges.

2.3 Characterizing the model as a two-equation system

We now show that any mixed-strategy equilibrium can be represented as a solution to a system of two equations in two unknowns. This result enables us to characterize the economy graphically and to obtain an intuition for the comparative statics that unfold when job search technology changes. To obtain the result, we first define the two type-specific job finding rates (f_L, f_H) , which will serve as the underlying system of variables in which the two equations are defined:

²³A trade-off between worker quality and the job filling rate has been documented in a growing empirical literature (Faberman and Menzio, 2018; Gulyas, 2025; Babalievsky et al., 2026). Davis, Faberman, and Haltiwanger (2013) also suggest that hiring standards are an instrument through which firms can steer vacancy filling rates.

²⁴See Appendix B.2 for formal definitions of an equilibrium and the special case of a mixed-strategy equilibrium.

$$f_i = \underbrace{\lambda}_{\text{Encounter rate}} \underbrace{p_i^{\text{app}}}_{\text{Application probability}} \underbrace{p_i^{\text{acc}}}_{\text{Acceptance probability}}, \quad p_i^{\text{acc}} = \begin{cases} (1 - p^{\text{sig}})(1 - R) + p^{\text{sig}}s, & i = H, \\ (1 - p^{\text{sig}})(1 - R) + p^{\text{sig}}(1 - s), & i = L. \end{cases} \quad (4)$$

Equation (4) formalizes that, in order for a worker to transition into a job, they have to encounter a job, apply to it, and be accepted. Acceptance entails being hired without a signal or through an H -signal, which is more likely to occur for H -type workers. As we will see below, f_L and f_H serve as sufficient statistics for all other endogenous objects of the model. This result enables us to write the model in terms of two equations in f_L and f_H .

The indifference curve. The first of the two equations comes from the first order condition of (3): In a mixed-strategy equilibrium, firms must be indifferent between accepting and rejecting workers without a signal. Using the fact that free entry implies $V = k$ in equilibrium, we thus obtain the following condition:

$$\underbrace{k}_{\text{Value of waiting}} = \underbrace{a_L J_L + (1 - a_L) J_H}_{\text{Value of hiring without signal}}. \quad (\text{IC})$$

A key insight that allows us to express this condition graphically is that all objects in this equation can be written as functions of the two job finding rates (f_L, f_H) and exogenous parameters; that is, $J_i = J_i(f_i)$ and $a_L = a_L(f_L, f_H)$. Appendix B.4 derives these functions in closed form. Intuitively, J_i depends negatively on f_i , since a higher job finding rate improves the outside option of the worker when bargaining and thus reduces the share of joint surplus that the firm can extract from a match. The share of L -type workers among applicants, a_L , rises in f_H and falls in f_L , since a larger job finding rate for a given worker type shifts the composition of the unemployed pool toward the remaining type.²⁵

Figure 7 plots the indifference curve for the parameters of our baseline calibration (see Section 2.5). In the plotted range, the curve is roughly horizontal, reflecting two opposing margins: as the curve moves right, f_L rises and the share of type- L applicants

²⁵Strictly speaking, this condition can be violated under a sufficiently large value of η_a , the elasticity of applications with respect to the value of applying. The reason is that, if η_a is large, a rise in f_L can in principle incentivize applications of L -types enough to make a_L increasing in f_L . However, we find that, in our calibration, a_L is decreasing in f_L and increasing in f_H in all relevant regions of the state space.




Figure 7: Indifference curve

Notes: The indifference curve is defined by all points (f_L, f_H) that solve equation (IC). $R = 0$ and $R = 1$ indicate regions of the state space in which firms optimally accept and reject no-signal encounters, respectively. Parameters are taken from Table 2.

falls, exerting upward pressure on the right-hand side of equation (IC). Simultaneously, the bargaining position of L -type workers improves, exerting downward pressure on the right-hand side of equation (IC). The two effects approximately cancel, leaving the right-hand side of equation (IC) unchanged as the curve moves to the right. Appendix B.5 discusses the shape of the indifference curve in more detail.

One instructive way to read the figure is as a map from the state of the economy, captured by (f_L, f_H) , to the firm's optimal hiring decision. Along the indifference curve, the value of hiring an applicant without a signal equals the value of waiting, so any $R \in [0, 1]$ is optimal. Consider instead a point above the curve. Holding f_L fixed, f_H is higher than the value required for indifference. The higher f_H improves the H -type worker's outside option and therefore lowers J_H . It also causes H -types to exit unemployment more quickly, shifting the applicant pool toward L -types and raising a_L . Consequently, at such points, $a_L J_L + (1 - a_L) J_H < k$, which makes rejecting an applicant without a signal uniquely optimal. For this reason, $R = 1$ is the optimal choice in states above the indifference curve. Conversely, below the curve, $R = 0$ is optimal.

The screening curve. The second condition that defines an equilibrium is an equation that captures the relationship of job finding rates between L -type and H -type workers under the assumed screening technology; we therefore refer to it as the screening curve.




Figure 8: Screening curve

Notes: The screening curve is defined by all points (f_L, f_H) that solve equation (SC). As indicated, points in the top-right direction correspond to a lower value of R . Parameters are taken from Table 2.

It is obtained by re-arranging the two expressions for f_L and f_H in (4):

$$f_H = \underbrace{\left(\frac{p_H^{\text{app}}}{p_L^{\text{app}}} \right) \cdot f_L}_{\text{Rescaled job finding rate of L-types}} + \underbrace{\lambda p_H^{\text{app}} p^{\text{sig}} (2s - 1)}_{\text{Screening advantage of H-types}} \quad (\text{SC})$$

Two margins introduce a wedge between f_L and f_H . First, the two types may apply at different rates, captured by the ratio of application probabilities. Second, type- H workers enjoy a screening advantage: since $s \geq 0.5$, a type- H worker gets hired with a higher probability than an L -type worker.

Just like before, the insight that simplifies the analysis of this equation is that $p_i^{\text{app}}, p^{\text{sig}}$, and λ can all be expressed as closed form functions of (f_L, f_H) and exogenous parameters, as we show in Appendix B.4. Intuitively, knowledge of job finding rates is sufficient to compute both the composition of applicants and their value to the firm, from which encounter rates, the signal probability, and application probabilities all follow.

Figure 8 plots the screening curve; that is, all combinations of (f_L, f_H) consistent with equation (SC). By equation (4), every point on the screening curve corresponds to a unique value of R : a lower value of R corresponds to higher job finding rates for both worker types, making the curve upward-sloping. Reducing R disproportionately raises the job finding rate of type- L workers, because their chances of being hired depend more




Figure 9: Equilibrium

Notes: The indifference curve defined by equation (IC) and the screening curve defined by equation (SC) are plotted together. Equilibrium job-finding rates lie at the intersection of the two curves. Parameters are taken from Table 2.

strongly on the likelihood that a firm admits workers without a signal.

Equilibrium. Since both (IC) and (SC) have to hold in a mixed-strategy equilibrium, the two type-specific job finding rates of the economy lie at the intersection of the two curves, as shown in Figure 9. All other equilibrium variables can be derived from (f_L, f_H) in closed form (see Appendix B.4). The equilibrium value of R is pinned down by the position at which the indifference curve intersects the screening curve.

2.4 Declining application costs through the lens of the model

Having cast the model as a system of two equations, we can now study how the economy responds to a decline in applications costs, modeled as an increase in γ_a . Figure 10 illustrates the resulting shift in equilibrium, using parameter values of the calibrated model, described in Section 2.5.

A decline in application costs leaves the indifference curve unchanged: the firm's decision between accepting and rejecting an applicant without a signal depends on the composition of the applicant pool and the value of filled and unfilled jobs, neither of which is affected by the cost of applying once we condition on type-specific job finding rates. The screening curve, however, shifts up: as applying becomes cheaper, the application probability of the H -type, p_H^{app} , rises, and the intercept term of the screening curve (SC)




Figure 10: Effect of an application cost shock

Notes: Red line: Indifference curve (IC) for the 1998-2000 and 2017-2019 steady state. Green solid line: Screening curve (SC) in 1998-2000 steady state. Green dashed line: Screening curve (SC) in 2017-2019 steady state (identical parameters except $\gamma'_a > \gamma_a$). Parameters are taken from Table 2. The blue circle labeled “Initial eqm.” and blue square labeled “New eqm.” indicate the respective equilibria; the blue triangle is the fixed- R counterfactual.

increases.

To build intuition for how the equilibrium shifts, we first hold R , the rejection probability of workers without a signal, at its initial value. The increase in applications raises job finding rates for both worker types, moving the economy from the initial equilibrium to the blue triangle in the upper right corner of Figure 10. To a first order, cheaper applications benefit both worker types by nearly the same amount, so this point lies above the initial screening curve.

The blue triangle, however, is not an equilibrium. In this state of the economy, the better outside option of workers improves their bargaining position. Therefore, the value of hiring a worker without a signal falls below the opportunity cost of hiring. As a result, firms now strictly prefer to reject such applicants in order to match with workers of higher expected productivity later on. R rises, which restores equilibrium by lowering the job finding rate of both worker types, but especially those of L -type workers. Therefore, after the shock, the economy moves to a new equilibrium indicated by the blue square. In the new equilibrium, the job finding rates of L -types are now strictly lower while those of H -types remain roughly unchanged.²⁶ The fall in application costs thus produces effects

²⁶While the equilibrium shift predicts an unambiguous fall in f_L , the effect on f_H can in principle be positive or negative depending on calibration choices. However, we find that, across calibrations, predicted

Parameter	Value	Target	Value
<i>Exogenously set / normalizations</i>			
Encounter productivity Z	$\approx 1.4 \cdot 10^3$	Normalization	—
Low-type productivity y_L	1.0	Normalization	—
Signal quality s	1.0	Normalization	—
Worker bargaining weight β	0.5	Convention (DMP)	—
Matching elasticity ω	0.3	Petrongolo and Pissarides (2001)	—
Discount rate r (monthly)	0.05/12	Convention	—
Destruction rate δ (monthly)	0.035	JOLTS	—
<i>Internally calibrated</i>			
H -type productivity y_H	10.1	Wage dispersion (CV)	0.47
Replacement factor $1 - \rho$	0.665	Replacement rate	0.4
Entry cost k	2.897	Job-finding rate	0.587
Ease of screening γ_s	3.300	Job-filling rate	1.75
Share of L -type workers α_L^{pop}	0.242	Unemployment duration (months)	3.0
Ease of applying γ_a	0.0039	Applications sent per month	2.8
Screening cost dispersion η_s	5.94	Screening costs / monthly output	0.048
Application cost dispersion η_a	0.391	Differential job finding elast.	0.763

Table 2: Parameter values and targets

Notes: Calibration table. “Parameter” indicates the name and symbol of the parameter. The first column labeled “Value” indicates the parameter value. “Target” indicates the targeted variable or the justification for exogenously set or normalized parameters. The second column labeled “Value” indicates the value of targeted variables. The model matches all targets to numerical precision.

that are in line with what we documented in Section 1: cheaper applications encourage firms to screen more thoroughly, leading to slower and more selective hiring.

2.5 Calibration

To compare the predictions of the model against the data, we next choose numerical values for the parameters of the model. Some parameter choices are normalizations, others we pick by convention. The remaining parameters target a set of salient labor market moments in the late 1990s, just before the advent of online job search. Table 2 summarizes all parameters values and their targets.

We begin by normalizing three parameters: Z , y_L , and s . We set the encounter productivity parameter Z to a value that guarantees a worker encounter rate of 1000. Next, we normalize the productivity of the low type y_L to one. Our final normalization is to set the signal quality s to one.²⁷

changes in f_H are typically very small when compared to the predicted change in f_L .

²⁷In Appendix C, we show that Z and s are normalizations in this model, because they can not separately

Next, we impose four parameters using conventional values from the literature: the worker bargaining weight $\beta = 0.5$, the elasticity of the encounter function $\omega = 0.3$ (Petrongolo and Pissarides, 2001), the discount rate r equal to 5% annually, and the job destruction rate $\delta = 0.035$, consistent with a long-run average of the separation rate according to data from JOLTS.

We use the remaining parameters to target moments of the model that reflect the US economy of the late 1990s. Since the distance between y_H and y_L governs the dispersion of wages in the model, we choose y_H to target a coefficient of variation of wages equal to 0.47, which we take from Hornstein et al. (2011). The replacement factor $1 - \rho$ targets a replacement rate of 40%, which is a standard value in the literature (Hall, 2005; Shimer, 2005).²⁸ Three parameters are chosen to target aggregate labor market series averaged from December 1998 to December 2000: the entry cost parameter k determines market tightness and is therefore chosen to target the average job finding rate. The screening cost parameter γ_s governs how easily firms can identify suitable applicants and is therefore chosen to target the average job filling rate. Lastly, we use the population share of L -type workers, α_L^{pop} , to target the average duration of unemployment. As we show in Appendix D.2, conditional on the job finding rate, unemployment duration is determined by the presence of worker heterogeneity, which is a function of the share of L -type workers in the population.

The three remaining parameters are the ease of applying, γ_a , and the application and screening elasticities, η_a and η_s . We choose the ease of applying, γ_a , to match an average monthly application rate of 2.8 among unemployed workers. This reflects the average rate of applications by unemployed workers in the EOPP.²⁹ Next, we turn to the application cost dispersion parameter η_a , which governs the sensitivity of type-specific application rates to the value of applying. We use this parameter to target the difference between the two type-specific elasticities of the job finding rate with respect to a change in unemployment caused by a productivity shock. In Appendix C, we show how we construct an empirical analog of this moment and obtain an empirical target of 0.763. Lastly, we set the screening cost dispersion η_s to target the cost of screening. The screening resources spent per month on an open position equal $k(r + \delta)\eta_s$ (derived in Appendix D.5) are directly proportional to η_s . In the 1982 National Survey of Employers (NES), we document that

be identified from γ_a and γ_s , respectively.

²⁸We define the replacement rate as the average production of unemployed workers relative to the average wage.

²⁹As stated above, the EOPP was conducted in 1980. However, to our knowledge, no application data exists from the period between 1980 and 2013, the first year of the SCE.

firms spend on average 8.3 personnel hours on recruiting, screening, and interviewing applicants per hire. This is equal to 4.8% of the monthly working hours of a typical full-time job. We therefore impose a value of η_s such that the screening resources spent per hire equal 4.8% of the monthly output of an average job. Appendix C provides additional details on the construction of the empirical targets we impose.

When conducting comparative statics below, we model online job search as a shock that decreases the cost of applications, increasing γ_a . We discipline the shock by imposing that it increases the average rate at which unemployed workers apply to jobs to 9.2, which corresponds to the value we observe in the SCE during 2013-22.

3 The effects of online job search in the model and data

Through the lens of the quantified model, we can now evaluate whether the application cost shock can explain the micro- and macro-level trends observed in Section 1. We begin by showing in Section 3.1 that the model produces shifts in aggregate variables that are in line with the data. We then show that the micro-level predictions of the model are also consistent with the evidence: first, we show in Section 3.2 that the model produces a positive shift in the composition of hires over time, in line with Section 1.3; next, in Section 3.3, we show that the dispersion in type-specific unemployment duration increases, as described in Section 1.4.

3.1 Aggregate labor market effects

Figure 11 plots time series of four aggregate labor market moments alongside the implied model changes in those moments. To sidestep the large role of the business cycle, we compare the data series across two points in time that feature a similar state of the business cycle; specifically, we compare our calibration period (December 1998 - December 2000) with a reference period of the last two pre-pandemic years (December 2017 - December 2019), two periods that feature average unemployment rates of 4.1% and 3.8%, respectively. Dashed lines indicate averages in the two time periods. We compare changes in the job finding rate, the average duration of unemployment, the job filling rate, and matching efficiency between model and data.

Job finding rate. Over the 20 year horizon during which online job search became ubiquitous, the job finding rate declined by 27% in the data. The model generates a decline




Figure 11: Comparison of aggregates across the two steady states

Source: Current Population Survey, Job Openings and Labor Turnover Survey, and authors' calculations.
Notes: Each panel plots a data series (solid black) alongside its averages in Dec. 1998 - Dec. 2000 and Dec. 2017 - Dec 2019 (dashed black lines for data, red for model). The 1998-2000 averages are targeted and 2017-19 averages are non-targeted by the model. Percentage changes between the two periods in the data and model are reported

in a similar range, at 22%. Two distinct forces contribute to this decline in the model. First, as illustrated in Figure 10, the job finding rate of *L*-type workers falls substantially, as firms become more selective by raising the rejection probability of workers without a signal. Second, the composition of the unemployment pool shifts toward *L*-type workers, who have lower job finding rates.

Unemployment duration. The average duration of unemployment rises by 68% in the model, closely matching the 71% increase in the data. Notably, in the model and in the data, the rise in unemployment duration substantially exceeds the decline in the job finding rate. This feature signals a departure from the homogeneous worker benchmark, since any search model without worker heterogeneity and constant exit hazards predicts that the two series should be inverses of one another. In fact, we derive in Appendix D.2 that the wedge between the average duration of unemployment and the inverse job finding rate captures the amount of worker heterogeneity in the model:

$$\underbrace{\text{unemployment duration}}_{\substack{+68\% \text{ in the model,} \\ +71\% \text{ in the data}}} = \underbrace{(\text{job finding rate})^{-1}}_{\substack{+29\% \text{ in the model,} \\ +37\% \text{ in the data}}} \cdot \underbrace{\left(1 + a_L^u(1 - a_L^u) \frac{(f_H - f_L)^2}{f_H f_L}\right)}_{+31\% \text{ in the model}}. \quad (5)$$

Equation (5) shows that the average duration of unemployment in our economy can be decomposed into the inverse job finding rate on the one hand and a measure of cross-type dispersion in job finding rates on the other. The model therefore reconciles a large increase in unemployment duration with a small decrease in the job finding rate because it predicts that, from the initial to the final steady state, job finding rates become more dispersed across the two types. The increase in heterogeneity among the unemployed that results from the shock contributes about half of the overall increase in unemployment duration.³⁰

Job filling rate. The job filling rate declines by 19% in the model and by 20% in the data. In the model, the job filling rate equals

$$\text{job filling rate} = \overbrace{q(a_L^u p_L^{\text{app}} + (1 - a_L^u) p_H^{\text{app}})}^{\text{rate at which applications are received}} \cdot \left((1 - p^{\text{sig}}) \underbrace{(1 - R)}_{\substack{\text{acceptance prob.} \\ \text{without signal}}} + p^{\text{sig}} \underbrace{[a_L(1 - s) + (1 - a_L)s]}_{\substack{\text{acceptance prob. with signal}}} \right), \quad (6)$$

where a_L^u is the share of L -type workers among the unemployed. Equation (6) shows that the job filling rate can be written as the product of two margins: the rate at which

³⁰Intuitively, heterogeneity in job finding rates among the unemployed matters for the duration of unemployment because of Jensen's inequality: as heterogeneity increases, the average of the inverse job finding rates ($\mathbb{E}_i [1/f_i]$), becomes larger than the inverse of the average job finding rate ($1/\mathbb{E}_i [f_i]$). We discuss this in more detail in Appendix A.7.

applications are received by firms, and the probability that an encounter converts into a hire. How the change in γ_a affects the job filling rate thus depends on how these two margins change.

Consider first the probability that an encounter converts into a hire, given by the second term in equation (6). Two forces push toward a decline in this term and therefore lower the job filling rate. First, firms reject a larger share of workers for whom they do not receive a signal (R increases). Second, the composition of the applicant pool shifts toward L -type workers (a_L increases), which reduces the probability of a favorable signal and resulting hire even when a signal is obtained.

Consider next the rate at which the firm receives applications. The shift in the composition of applicants induced by γ_a causes a decline in the firm's return to purchasing a signal, since the new applicant pool has a higher share of L -types and a lower variance of type-specific job values. Free entry implies that the firm must receive applications more frequently to keep the value of a vacancy equal to the sunk entry cost. The rate at which the firm receives applications thus rises after the shock and exerts upward pressure on the job filling rate. This channel is less powerful in our parametrization, leading to an overall decrease in the job filling rate.³¹

The model thus rationalizes the simultaneous decline of job finding and filling rates, which reflects a two-sided slowdown of hiring rates in the market. Our final observation below is therefore that the model delivers long-run changes in a typical measure of matching efficiency that mirror those in the data.

Matching efficiency. Matching efficiency is a summary measure that captures how effectively the labor market produces matches between job seekers and prospective employers. One common definition of matching efficiency is the ratio of the number of hires observed and the number of hires produced by an assumed matching function: $\text{hires} / \sqrt{\text{unemployment} \cdot \text{vacancies}}$. This measure declines by 21% in the model, compared to 25% in the data. In the model, the application cost shock lowers the fraction of contacts that result in a hire, which leads to fewer hires per unit of search activity. The model therefore reproduces the long term decline in matching efficiency observed in the data.

³¹Appendix D.1 provides a formal derivation of the free entry condition which relates the signal probability p^{sig} with the rate at which applications are received.

3.2 Positive selection of hires

We next test whether the model can account for the compositional changes in hires that we documented in Section 1.3. There, we showed that, relative to the 1990s, the composition of hires has increasingly shifted towards workers with higher pre-separation wages and relatively shorter unemployment durations. To see whether the model is consistent with these trends, we construct both pre-separation wage ranks and unemployment duration ranks of newly hired workers in the model. We do this in a way that is designed to mirror the empirical approach from Section 1.3.

We first construct the pre-separation rank of hires by using the law of motion of the employed to project forward a cohort of workers employed in a reference month 0. Since the employment state of each worker follows a two-state Markov process, it is possible to use the Markov transition matrices of each worker to compute the average rank of workers among that cohort who are being hired 12 months later. Appendix D.3 provides details on the construction.

We then construct a measure of the unemployment duration rank that is analogous to the one we use in Section 1.3. The unemployment duration rank in the model is the average rank that new hires occupy in the distribution of unemployment duration at the time of their hire. As we show in Appendix D.4, this object admits a closed-form expression:

$$\text{unemployment duration rank of new hires} = \frac{1}{2} \left(1 - \frac{a_L^u(1 - a_L^u)(f_H - f_L)^2}{(f_L + f_H)(a_L^u f_L + (1 - a_L^u)f_H)} \right).$$

This expression shows that diverging job finding rates among the two worker types exert downward pressure on the unemployment duration rank of new hires: as low types remain persistently unemployed, the composition of hires shifts toward workers with relatively lower duration.

Having defined the model counterparts from the series in Section 1.3, we can now compare how the model performs in replicating the observed shifts in the composition of hires. Figure 12 shows both rank measures from Figure 3 side by side with the model-implied changes. The model implies an increase in the pre-separation wage rank of new hires by about 1 percentage point, somewhat below the roughly 2 percentage point increase in the data. The model reflects movements in the unemployment duration rank of new hires more accurately: it implies a decline by 1.3 percentage points, close to the




Figure 12: Pre-separation and unemployment duration ranks in model and data

Source: Current Population Survey and authors' calculations. *Notes:* Each panel plots a data series (solid black) alongside its averages in Dec. 1998 - Dec. 2000 and Dec. 2017 - Dec. 2019 (dashed black lines for data, red for model). Percentage changes between the two periods in the data and model are reported.

decline of 1.4 percentage points in the data.

The model thus captures that the quality of hires improves in a way that is consistent with the data: since firms become more selective in who they hire, new hires in the post-internet steady state feature both higher pre-separation wage ranks and shorter relative unemployment spells compared to the pre-internet era. This increase in selectivity that is visible in the composition of hires has implications for the relative job finding prospects of different worker types, which we study next.

3.3 Heterogeneous trends in unemployment duration

In Section 1.4, we documented that the sustained increase in unemployment duration was concentrated among a subset of workers. Table 3 shows that the model is capable of reproducing this pattern. Whereas the unemployment duration of *H*-type workers changes by less than one percent, the average unemployment duration of *L*-types increases substantially, by 57%.

To compare these magnitudes with the data, we again utilize the NLSY where, like in Section 1.4, we classify workers by the lifetime average duration of their unemployment spells. To obtain a comparison that is model-consistent, rather than classifying by

Type	Level (1998-00)	Level (2017-19)	Change	Change (NLSY)
<i>L</i>	4.495	7.069	+57.3%	+73.4%
<i>H</i>	0.911	0.904	-0.8%	+12.6%

Table 3: Type-specific unemployment duration across steady states in model and data

Notes: Each row corresponds to a variable. “Level” columns correspond to model values of the corresponding variable in the indicated steady state. The “Change” column denotes the percentage change between both steady states. The last column labeled “Change (NLSY)” shows the the type-specific changes observed in the data, based on the classification we construct in the NLSY, as described in the main text.

whether a worker is above or below the median of the distribution of lifetime unemployment duration, we instead divide the population into quartiles and classify the quartile of workers with the longest lifetime average unemployment duration as *L* and workers who belong to one of the remaining three quartiles as *H*.³² Since age is an important predictor of type-specific unemployment duration, we focus on a cross-cohort comparison and compare the average duration of each cohort’s first 17 observed years by type.³³ For each type we thus compare the average unemployment duration in the NLSY 97 cohort from 2005 to 2021 with the average in the NLSY 79 cohort from 1982 to 1998. Table 3 reports this change. While *L* types saw substantially longer unemployment spells, the duration of *H* types changed relatively little across the two cohorts. The percentage increases of 73.4% for *L*-types and 12.6% for *H*-types are a bit larger but broadly consistent with the model counterparts of 57.3% and -0.8%, respectively.

One implication of this asymmetric decline in job finding rates is that they should fall more for workers with longer unemployment spells; in other words, the aggregate job finding rate should become more duration dependent over time. Figure 13 shows that this is indeed the case both in the model and in the data. The left-hand side panel shows the log probability of a UE transition in the model for workers who have been unemployed for different amounts of time, both before and after the shock. The right-hand side panel shows the same statistic using CPS data: we compute, for workers with different elapsed unemployment spell lengths, the probability with which they transition into employment in the next month. The two lines correspond to data from 1998–2000 and 2015–2019, respectively. Both the model and the data show a steeper profile in the later period (red line), reflecting an increase in duration dependence and a fall in job finding

³²This reflects that α_L^{pop} , the share of *L*-type workers in the population, is 0.24 in our calibration.

³³We choose a period of 17 years because this is the entire time span for which we observe workers in the NLSY 97. This approach implicitly conditions on age.




Figure 13: Job finding probability by unemployment duration, model vs. data

Source: Current Population Survey. *Notes:* Both panels plot the log probability of becoming employed in the following month by elapsed unemployment duration. The left panel compares the model's steady states before (blue) and after (red) the decline in application costs. The right panel compares CPS estimates for 1998-2000 (blue) and 2015-2019 (red). Shaded bands indicate 95% confidence intervals.

rates that is concentrated among the long-term unemployed.

Taking stock. Our results show that a change in job search technology that facilitates applications changes two important margins that are relevant from a macroeconomic perspective. First, the shock changes who gets hired: hires shift towards workers with higher productivity, which explains the pronounced and concentrated rise in unemployment duration observed in the data since the 1990s. Second, more selective hiring slows matching rates, explaining the observed decline in job finding and filling rates, as well as matching efficiency, since the 1990s.

4 The impact of artificial intelligence on matching markets

The model successfully captures the central mechanism through which a decline in application costs affects the labor market: it incentivizes screening. However, application costs are just one of several changes in job search technology that affect firms' incentives to screen. Another important margin is the cost and quality of information. The idea that changes in information frictions may affect labor market matching has recently been em-

phasized in the context of artificial intelligence (AI) and its deployment in job search: a growing empirical literature (Cowgill et al., 2024; Galdin and Silbert, 2025) finds that job application materials are becoming less informative about their senders due to the use of large language models (LLMs) in writing. Our model offers a laboratory to explore the macroeconomic effects of such an innovation. In this final section, we thus apply the model to understand the macroeconomic effects of the use of LLMs in job search. We begin by studying the impacts of changes in the cost and quality of screening (Section 4.1) before offering a back-of-the-envelope quantification of the LLM shock and its effects (Section 4.2).

4.1 Equivalence of application and information shocks

How do changes to the cost of screening (γ_s) and quality of screening (s) affect the labor market? The model delivers a surprisingly simple answer: Both incentivize screening in the same way that a shock to the quality of applying does, with identical downstream effects. Put differently, shocks to γ_a , γ_s , and s , are equivalent in our model. To see why this is the case, it is useful to re-write the screening curve as a function of a sufficient statistic that depends on the three parameters, as we do in the following proposition:

Proposition 1 (Sufficient statistic for screening). *The screening curve (SC) can be written as*

$$f_H = M_1(\vec{f}) \cdot f_L + X(\gamma_a, \gamma_s, s)^{\frac{1}{1+\eta_a-\omega}} \cdot M_2(\vec{f}).$$

where

$$X(\gamma_a, \gamma_s, s) = \gamma_a \gamma_s^{1+\eta_a} (2s-1)^{(1+\eta_s)(1+\eta_a)}.$$

and $M_1(\vec{f})$, $M_2(\vec{f})$ do not depend on (γ_a, γ_s, s) .

Proof. See Appendix D.6. □

Proposition 1 implies that an increase in γ_a shifts the screening curve in exactly the same way as an equal increase in $\gamma_s^{1+\eta_a}$ or $(2s-1)^{(1+\eta_s)(1+\eta_a)}$. The indifference curve, meanwhile, does not depend on either parameter. For any shock to γ_a , there are thus corresponding shocks to γ_s and s that introduce the same movement in almost all aggregate variables. We formalize this insight in Proposition 2:

Proposition 2 (Equivalence of shocks). *There exists $\varepsilon > 0$ such that, for any pair (γ_a, γ'_a) with associated mixed-strategy equilibria (E, E') that feature $(f'_L, f'_H) \in B_\varepsilon((f_L, f_H))$, the value*

$$s' = \frac{1}{2} \left[1 + (2s - 1) \left(\frac{\gamma'_a}{\gamma_a} \right)^{\frac{1}{(1+\eta_s)(1+\eta_a)}} \right]$$

and the value

$$\gamma'_s = \gamma_s \left(\frac{\gamma'_a}{\gamma_a} \right)^{\frac{1}{1+\eta_a}}$$

are associated with mixed-strategy equilibria E'_s and E'_{γ_s} . The three equilibria $(E', E'_s, E'_{\gamma_s})$ feature the same type-specific job finding rates, distribution of unemployment duration, type-specific unemployment stocks, vacancy stocks, job filling rates, and matching efficiency. The equilibria E'_s and E'_{γ_s} feature identical type-specific application rates. The equilibria E'_{γ_a} and E'_{γ_s} feature identical rates of signal extraction.

Proof. See Appendix D.6. □

Proposition 2 states that shocks to γ_a , γ_s and s induce the same changes in aggregate variables (though they may have different implications for application and screening rates).³⁴ Intuitively, this result arises because positive changes to all three variables shift the screening curve upward, which induces firms to become more selective (increasing R) and shifts the equilibrium job finding rates along the indifference curve to the left.

The model thus reveals that the effects of LLMs in job search depend on the relative strength of three competing effects. First, LLMs may make it easier for workers to apply (increasing γ_a). Second, LLMs may make it easier for firms to scan applicants at scale and therefore lower the cost of extracting signals about applicants (increasing γ_s). Lastly, as emphasized in Cowgill et al. (2024) and Galdin and Silbert (2025), the new technology may lower the quality of information firms can extract from application portfolios, since lower-quality applicants can use LLMs to improve and customize their application materials (decreasing s). Proposition 2 makes clear that this last effect has the opposite sign compared to the first two.

³⁴The central qualification of Proposition 2 is that these changes have to be small enough ($< \varepsilon$) that the model does not leave the mixed-strategy region. We find that, in practice, ε can be chosen to be large enough to guarantee equivalence of any γ_a , γ_s and s shocks in realistic ranges.

A serious quantification of each channel would require reliable data on the effect that LLMs have on the cost of applying, the cost of screening, and the quality of information. Unfortunately, the available data on the effects of LLMs along these margins is limited. Therefore, instead of a serious quantification, we use the last part of this paper to make a best-effort attempt to attach numbers to each of the three channels while noting that our estimates are a back-of-the-envelope calculation.

4.2 Assessing the effects of LLMs on matching

Our quantification is based on three ideas: first, we use intertemporal changes in application rates to identify changes in γ_a , second, we use intertemporal changes in the rate of interviews to identify changes in γ_s , and third, we use experimental evidence from [Cowgill et al. \(2024\)](#) to identify changes in s .³⁵ We discuss each in turn.

First, to quantify the shock to γ_a , we rely on reporting from Greenhouse, a large job search platform, that the number of applications per job posting approximately doubled from the release of ChatGPT in November 2022 until the end of 2025.³⁶ Motivated by this report, we assume the change in γ_a necessary to double the number of applications per hire: $\gamma_a^{\text{AI}} / \gamma_a^{\text{noAI}} = 2.37$.

Next, we use the rate of screening to discipline the change in γ_s . While the screening rate is a model-specific object that is not directly observed in the data, we make the identifying assumption that the screening rate (i.e., the rate of signaling events) is proportional to the rate of interviews observed in the data.³⁷ Under this assumption, the change in the rate of interviews per hire equals the change in the rate of signals per hire. According to Ashby, a human resource company, the rate of interviews per hire increased by approximately 31% between Q3 2022 and Q1 2026.³⁸ We find that the model generates such a change when $\gamma_s^{\text{AI}} / \gamma_s^{\text{noAI}} = 2.12$.

Lastly, we quantify the change in s . We rely on experimental evidence by [Cowgill et al. \(2024\)](#) who collect experimental evidence on the change in posterior that recipients of applications experience when the senders have access to LLMs. We relegate details to

³⁵By Proposition 2, only application rates and screening rates can be used to separately identify the underlying contribution of γ_a , γ_s and s to changes in the aggregate environment. Our strategy is therefore to use application rates and screening rates as identification for two of the three shocks and use direct experimental evidence to identify the third.

³⁶See Business Insider, “[How Tech Broke the Job Market](#)”, accessed March 19, 2026.

³⁷This would be true in a re-interpretation of the model that casts interview costs d as distributed with cdf $\tilde{G}(d) = \tau^{-1}G(d/\tau)$ and views signals as occurring with probability τ at each interview. This setup recovers all equations that define the model.

³⁸See Ashby, “[Recruiter Productivity 2026 Talent Trends Report](#)”, accessed August 24, 2026.




Figure 14: Shifts in the screening curve resulting from AI

Notes: We label the re-calibrated 2019 steady state “Pre AI”. The green solid line displays the initial screening curve (SC). The blue dashed line, the green dash-dotted line and the orange dotted line display successive shifts as γ_a , then γ_s , and finally s are set to their post-AI values.

Appendix C.3. Mapping the estimates from Cowgill et al. (2024) to our model, we find that the signal accuracy $2s - 1$ declines by 6.5%.

To analyze the macroeconomic effects of these changes in search technology, we then re-calibrate the model to the 2017-19 period.³⁹

Figure 14 shows how easier applications (γ_a^{AI}), easier screening (γ_s^{AI}) and a worse signal quality (s^{AI}) shift the screening curve. Because AI makes applying to jobs cheaper, the screening curve shifts up and the equilibrium shifts to the left, with negative effects on the job finding rates of L -type workers but small (mildly positive) effects on the job finding rates of H -type workers. The lower cost of screening has similar effects and shifts the curve further to the left. The negative shock to the quality of the signal offsets some, but not all of the other two effects. The final steady state therefore lies to the left of the initial steady state.

Table 4 details the combined effects of these changes on aggregate variables and compares the implied changes in aggregate labor market variables with the change between the pre-pandemic steady state (December 2017 to December 2019) and the post-AI labor market (October 2024 to December 2025). We find that changes in job search technology fueled by AI explain about 63% of the decline in the job finding rate (13.9% versus 22.2%),

³⁹We use the same aggregate targets as in the 1998-00 calibration but target the period of December 2017 to December 2019. We keep η_a and η_s at their initially estimated values and drop their respective targets.

Variable	γ_a^{AI}	$+\gamma_s^{\text{AI}}$	$+s^{\text{AI}}$	Model change (2019–2025)	Data change (2019–2025)
Job finding rate	-10.3%	-9.2%	+5.8%	-13.9%	-22.2%
Unemployment duration	+25.8%	+23.0%	-11.4%	+37.1%	+16.9%
Job filling rate	-5.8%	-1.4%	+0.3%	-6.8%	-8.4%

Table 4: Labor market impacts of the LLM technology

Notes: Each row corresponds to a model variable. The “ γ_a^{AI} ” column indicates the percentage change of the variable when $\gamma_a = \gamma_a^{\text{AI}}$ but all other variables are as in the 2019 steady state. The columns labeled “ $+\gamma_s^{\text{AI}}$ ” and “ $+s^{\text{AI}}$ ” indicate the residual percentage changes of each variable when, in addition, γ_s^{AI} and $s = s^{\text{AI}}$ are set to their post-AI value, respectively. Changes in the “Data” column refer to changes in averages between (December 2017 - December 2019) and (October 2024 - December 2025).

overexplain the rise in unemployment duration (37.1% versus 16.9%), and explain about 81% of the decline in the job filling rate (6.8% versus 8.4%).

5 Conclusion

Improvements in job search technology can account for long-term labor market trends observed since the 1990s, in particular the uneven lengthening of unemployment spells and the fall in matching efficiency. Our framework helps understand the role of internet job search in amplifying search frictions and lets us interpret recent changes brought about by AI. By continuing to make applications easier and screening less costly, AI has the potential to further intensify congestion in hiring.

The broader lesson of our paper is that lower search frictions brought about by recent improvements in search technology may, paradoxically, worsen congestion frictions and make labor market access more unequal. In markets where firms select their preferred type of worker among a large pool of applicants, technologies that facilitate contacts can generate more selectivity, fewer hires, and more unequal job finding prospects. Identifying policies and market designs that preserve the benefits of easier search without amplifying these congestion forces is an important direction for future work.

References

- Ahn, H. J. (2023). The role of observed and unobserved heterogeneity in the duration of unemployment. *Journal of Applied Econometrics* 38(1), 3–23.
- Ahn, H. J. and J. D. Hamilton (2020). Heterogeneity and unemployment dynamics. *Journal of Business & Economic Statistics* 38(3), 554–569.
- Auster, S., P. Gottardi, and R. Wolthoff (2025). Simultaneous search and adverse selection. *Review of Economic Studies*, rdaf014.
- Babalievsky, F., H. R. Hyatt, and L. F. Warren (2026). Hiring with a quantity-skill tradeoff. Work in progress.
- Bagga, S., L. F. Mann, A. Şahin, and G. L. Violante (2025). Job amenity shocks and labor reallocation. Technical report, National Bureau of Economic Research.
- Baley, I., A. Figueiredo, and R. Ulbricht (2022). Mismatch cycles. *Journal of Political Economy* 130(11), 2943–2984.
- Barnichon, R. (2010). Building a composite help-wanted index. *Economics Letters* 109(3), 175–178.
- Bhuller, M., D. Ferraro, A. R. Kostøl, and T. C. Vigtel (2023). The internet, search frictions and aggregate unemployment. Technical report, National Bureau of Economic Research.
- Birinci, S., K. See, and S. L. Wee (2025). Job applications and labour market flows. *Review of Economic Studies* 92(3), 1438–1496.
- Blanchard, O. J., P. Diamond, R. E. Hall, and K. Murphy (1990). The cyclical behavior of the gross flows of us workers. *Brookings papers on economic activity* 1990(2), 85–155.
- Bradley, J. (2025). Worker-firm screening and the business cycle. *Review of Economic Dynamics* 57, 101272.
- Bradley, J. and L. Mann (2024). Learning about labor markets. *Journal of Monetary Economics* 148, 103612.
- Cai, X., P. Gautier, and R. Wolthoff (2025). Search, screening, and sorting. *American Economic Journal: Macroeconomics* 17(3), 205–236.

- Cowgill, B., P. Hernandez-Lagos, and N. L. Wright (2024). Does ai cheapen talk? theory and evidence from global entrepreneurship and hiring. *Theory and Evidence From Global Entrepreneurship and Hiring* (July 26, 2024). *Columbia Business School Research Paper* (4702114).
- Cui, J., G. Dias, and J. Ye (2025). Signaling in the age of ai: Evidence from cover letters. *arXiv preprint arXiv:2509.25054*.
- Davis, S. J., R. J. Faberman, and J. C. Haltiwanger (2013). The establishment-level behavior of vacancies and hiring. *The Quarterly Journal of Economics* 128(2), 581–622.
- Dhillon, P., Y. Gao, T. Lu, and Y. Zhang (2025). From signal to noise: How widespread llm usage transforms evaluator behavior in credit screening. *Available at SSRN* 5616650.
- Diamond, P. A. and A. Şahin (2016). Disaggregating the matching function. Technical report, National Bureau of Economic Research.
- Engbom, N. (2021). Contagious unemployment. Technical report, National Bureau of Economic Research.
- Faberman, R. J. and G. Menzio (2018). Evidence on the relationship between recruiting and the starting wage. *Labour Economics* 50, 67–79.
- Fujita, S., G. Moscarini, and F. Postel-Vinay (2024). Measuring employer-to-employer reallocation. *American Economic Journal: Macroeconomics* 16(3), 1–51.
- Galdin, A. and J. Silbert (2025). Making talk cheap: Generative ai and labor market signaling. *arXiv preprint arXiv:2511.08785*.
- Gans, J. S. (2024). How will generative ai impact communication? *Economics Letters* 242, 111872.
- Gulyas, A. (2025). The puzzling labor market sorting pattern in expanding and contracting firms. Conditionally accepted at the Review of Economic Studies.
- Hall, R. E. (2005). Employment fluctuations with equilibrium wage stickiness. *American Economic Review* 95(1), 50–65.
- Hornstein, A., P. Krusell, and G. L. Violante (2011). Frictional wage dispersion in search models: A quantitative assessment. *American Economic Review* 101(7), 2873–2898.

- Jarosch, G. and L. Pilossoph (2019). Statistical discrimination and duration dependence in the job finding rate. *The Review of Economic Studies* 86(4), 1631–1665.
- Kroft, K., F. Lange, M. J. Notowidigdo, and L. F. Katz (2016). Long-term unemployment and the great recession: the role of composition, duration dependence, and nonparticipation. *Journal of Labor Economics* 34(S1), S7–S54.
- Landais, C., P. Michaillat, and E. Saez (2018). A macroeconomic approach to optimal unemployment insurance: Theory. *American Economic Journal: Economic Policy* 10(2), 152–181.
- Mann, L. (2026). Labor market selection and the dynamics of a recovery.
- Martellini, P. and G. Menzio (2020). Declining search frictions, unemployment, and growth. *Journal of Political Economy* 128(12), 4387–4437.
- Mueller, A. I. (2017). Separations, sorting, and cyclical unemployment. *American Economic Review* 107(7), 2081–2107.
- Mueller, A. I., J. Spinnewijn, and G. Topa (2021). Job seekers perceptions and employment prospects: Heterogeneity, duration dependence, and bias. *American Economic Review* 111(1), 324–363.
- Mukoyama, T. and A. Şahin (2009). Why did the average duration of unemployment become so much longer? *Journal of Monetary Economics* 56(2), 200–209.
- Petrongolo, B. and C. A. Pissarides (2001). Looking into the black box: A survey of the matching function. *Journal of Economic literature* 39(2), 390–431.
- Petrosky-Nadeau, N. and L. Zhang (2021). Unemployment crises. *Journal of Monetary Economics* 117, 335–353.
- Pries, M. J. and R. Rogerson (2022). Declining worker turnover: The role of short-duration employment spells. *American Economic Journal: Macroeconomics* 14(1), 260–300.
- Shimer, R. (2005). The cyclical behavior of equilibrium unemployment and vacancies. *American Economic Review* 95(1), 25–49.
- Shimer, R. (2012). Reassessing the ins and outs of unemployment. *Review of Economic Dynamics* 15(2), 127–148.

- Stigler, G. J. (1961). The economics of information. *Journal of political economy* 69(3), 213–225.
- Stigler, G. J. (1962). Information in the labor market. *Journal of political economy* 70(5, Part 2), 94–105.
- U.S. Census Bureau (2019, October). Current Population Survey: Design and methodology. Technical Paper 77, U.S. Census Bureau.
- Wiles, E. and J. J. Horton (2025). Generative ai and labor market matching efficiency. *Available at SSRN 5187344*.
- Wolthoff, R. (2018). Applications and interviews: Firms recruiting decisions in a frictional labour market. *The Review of Economic Studies* 85(2), 1314–1351.

A Data appendix

A.1 Data on job applications

This appendix describes the two data sources underlying the right panel of Figure 1: the SCE and EOPP surveys.

Survey of Consumer Expectations. For 2013-22, we measure applications using the SCE Job Search Supplement, fielded annually in October. Respondents report the number of applications submitted by answering the question *“How many potential employers, if any, did you apply to for employment within the last 4 weeks? Please include all applications made in person, online, or through other direct methods. Do not include inquiries that did not lead to a job application.”* We restrict the sample to respondents who were unemployed in the month preceding the interview and weight observations using the survey weights. To trace the evolution of the distribution, we pool survey waves into three periods: 2013-16, 2017-19, and 2021-22.

Employment Opportunity Pilot Project. The EOPP was a household survey fielded in 1980. It included a “Looking-for-Work” record which asked respondents, for each spell of job search since January 1979 and each search method they used, *“How many applications did you fill out as a result of [search method]?”* Applications are therefore recorded per search method and search spell, whereas the SCE asks about the four weeks before the interview. To make the EOPP comparable with the SCE, we first sum applications across all search methods within each spell. We then prorate each spell’s total applications by the share of the spell falling within the four weeks preceding the interview date, assuming a constant application rate over the spell. For example, a spell that falls entirely within the four-week window contributes all of its applications, a spell that overlaps the window for a third of its length contributes a third, and a spell that ended before the window contributes zero applications. Finally, we sum applications across search methods within a spell, and across overlapping spells within each respondent. If the application count is missing for any overlapping spell, we set the respondent’s four-week measure to missing. All estimates use individual survey weights.

A.2 Job filling rate and recruitment of HR specialists

This appendix describes the sources underlying the construction of the job filling rate and firms' demand for recruiting and screening in Figure 2 and discusses their robustness to alternative measurement.

Construction of the job filling rate. The job filling rate is measured as the ratio of the monthly flows of hires to the end-of-month stock of vacancies. We take quarterly averages of monthly data. Vacancies are JOLTS job openings from the fourth quarter of 2000 onward, where a job opening is a specific position that could start within thirty days and for which the employer is actively recruiting outside the establishment. Before the JOLTS became available, we use the composite help-wanted vacancy rate of [Petrosky-Nadeau and Zhang \(2021\)](#), which we convert into a vacancy stock using CPS employment: $V_t = [v_t / (1 - v_t)] E_t$ where V and E denote vacancy and employment stocks, and v denotes the vacancy rate, which from January 1995 to November 2000 captures the composite print and online help-wanted index from [Barnichon \(2010\)](#) and, before January 1995, equals the help-wanted advertising index from the Conference Board.

We measure hires from the CPS as a sum of hires from unemployment, hires from non-participation, and hires from employment. The first two are constructed from CPS gross flows as the product of transition probabilities and relevant population stocks. Hires from employment are constructed by multiplying the employment stock with the Employer-to-Employer transition probability constructed from [Fujita, Moscarini, and Postel-Vinay \(2024\)](#) post 1995, and [Diamond and Şahin \(2016\)](#) and [Blanchard, Diamond, Hall, and Murphy \(1990\)](#) prior to 1995. After constructing total hires in the CPS, we adjust the hires following the methodology of [Bagga, Mann, Şahin, and Violante \(2025\)](#): over the period when the JOLTS and CPS overlap from 2001 onward, we rescale CPS hires to preserve the level of hires while matching JOLTS fluctuations in hires. We combine the adjusted hires series with the unadjusted CPS hires series before 2001 when JOLTS was not available. This yields a series that has a longer frequency, and is consistent in levels with the CPS data and higher-frequency movements with the JOLTS data.

Figure A.1 shows that our qualitative results do not hinge on this measurement choice. The figure plots three measures of job filling rate: our benchmark measure using adjusted CPS hires divided by vacancies (red solid), unadjusted CPS hires divided by vacancies (green dash-dot), and JOLTS hires divided by JOLTS vacancies (blue dash). The unadjusted and adjusted CPS-based series track each other closely throughout the sample,




Figure A.1: Alternative measures of job filling rate

Source: Current Population Survey, Job Openings and Labor Turnover Survey, and [Petrosky-Nadeau and Zhang \(2021\)](#). Notes: The figure compares three measures of monthly hires divided by the stock of vacancies. The benchmark (red solid) adjusts CPS hires to preserve their average level while matching JOLTS fluctuations following the methodology of [Bagga, Mann, Şahin, and Violante \(2025\)](#). The alternatives use unadjusted CPS hires (green dash-dot) and JOLTS hires (blue dashed).

while the JOLTS-based series lies below both, but exhibits the same cyclical properties and the same downward trend post 2000.

Construction of HR-specialist vacancy and employment shares. We proxy for firms' demand for labor tasked with recruiting and screening by tracking demand for an occupation called Human Resources Specialists. In the 2010 and 2018 Standard Occupational Classification (SOC), Human Resources Specialists (13-1071) are defined as workers who *"perform activities in the human resource area. Includes employment specialists who screen, recruit, interview, and place workers."* In the 2000 SOC, the same code denotes "Employment, Recruitment, and Placement Specialists," defined as workers who *"recruit and place workers."* In all three classifications, the occupation excludes "Compensation, Benefits, and Job Analysis Specialists" and "Training and Development Specialists", who perform HR functions unrelated to hiring.

The right panel of Figure [A.2](#) plots employment in SOC 13-1071 as a share of em-




Figure A.2: Screening personnel: vacancy and employment share

Source: Lightcast, and Occupational Employment and Wage Statistics. *Notes:* The left panel plots HR specialists' (SOC 13-1071) share of online job postings and the right panel plots their share of employment. In both panels, blue solid lines use business and financial operations occupations (SOC 13) as the denominator (left axis), and red dashed lines use all occupations (right axis).

ployment in SOC 13 (Business and Financial Operations Occupations), using the May estimates of the Occupational Employment and Wage Statistics (left axis). As with job ads, we normalize by the two-digit family of HR specialists because employment has become concentrated in larger firms over this period, which mechanically raises the demand for all business operations staff. When measured as a fraction of all occupations, the share rises more steeply (right axis). Both employment series are interrupted by the 2010 revision of the SOC, and the OEWS reporting only hybrid codes in 2010-11 making the occupation unobservable. The figure shows that just like the increase recorded in the corresponding job ads series, employment of HR specialists rose from about 6.1% of business and financial operations occupations in 2012 to about 8.9% by 2024.

A.3 Selection of hires and unemployment duration

Construction of pre-displacement wage and unemployment duration rank of hires. We construct both measures in Figure 3 using the Basic Monthly Surveys of the CPS for individuals aged 16 to 64. The CPS interviews households for four consecutive months




Figure A.3: Wage rank of hires and separators

Source: Current Population Survey. *Notes:* The figure compares the mean pre-displacement wage ranks of hires (red solid) and separators (blue dashed). Ranks are computed within each year after residualizing log hourly wages with respect to age, education, and occupation. Both series are centered three-year moving averages, normalized to zero in 1985.

(interview months 1-4), leaves them out of the sample for eight months, and then interviews them again for another four months (interview months 5-8). We measure wages in interview month 4, the first outgoing rotation group interview. Our measure of wages combines reported hourly wages for workers paid by the hour, and the ratio of weekly earnings by weekly hours for other workers. We trim the individual wage distribution at the 1st and 99th percentiles and deflate wages using the Consumer Price Index. We then regress log hourly wages on age, three education groups, two-digit time-consistent occupation groups, and year fixed effects. Finally, we rank individuals by their residual wage within the year in which the wage is observed. Ranks are expressed on the unit interval, with a value of 0.5 corresponding to the median earner in that year.

We follow workers with a valid wage observation in interview month 4 through interview months 5 to 8. We classify a worker as one hired from non-employment in year t if they are employed in interview month 8 and were not employed in at least one of interview months 5 through 7. The left panel of Figure 3 plots the annual weighted mean rank of these workers, normalized to zero in 1985, the first year shown in the figure.

As our wage rank measure captures only those workers who were employed approx-

imately a year before being hired, workers with longer spells of non-employment may be underrepresented. We therefore complement this measure with a measure based on unemployment duration that does not condition on prior employment status. To construct this measure, we utilize measures of unemployment duration in the monthly CPS. Before 1994, the CPS asked unemployed workers “*How many weeks have you been looking for work?*” In 1994, the question changed to “*As of the end of LAST WEEK, how long had you been looking for work?*” (U.S. Census Bureau, 2019). Based off these questions, the CPS reports unemployment duration of unemployed workers in every month. We consider all working-age unemployed respondents in the monthly CPS each year between 1976 to 2023. We residualize the log of their elapsed unemployment duration with respect to age, education, and occupation, and then rank the unemployed workers each year by their residual duration. We classify a worker as a hire if they are unemployed in month t and employed in month $t + 1$. The right panel of Figure 3 reports the annual weighted average duration rank among hires, normalized to zero in 1976.

A.4 Rising duration: alternative explanations and robustness

Our benchmark measure of unemployment duration in Figure 4 is the BLS-published measure of the average time spent in unemployment among unemployed persons in a month. The upper left panel of Figure A.4 shows that the aggregate increase reflects a broader shift in the duration distribution. Using CPS micro data, we plot the shares of unemployed with elapsed duration less than one month and more than six months, alongside mean elapsed duration. We exclude temporary layoffs from unemployment. The short duration share trends downward over the sample, falling from nearly 45% in the late 1970s to about 30% by the end of the sample. The long-duration shares offset this decrease by moving in the opposite direction. The figure illustrates that the rise in average unemployment duration documented in Figure 4 is driven by a thickening of the right tail of the duration distribution, that is, a shift away from short ongoing spells toward longer ones.

We discuss three explanations distinct from our own that could potentially account for the rise in mean unemployment duration. First, over the last 50 years, the composition of labor force has shifted toward older and more educated workers, and female labor force participation has increased. If these groups experience longer than average unemployment spells, then changes in the composition of the unemployed could drive the aggregate increase in mean duration. To control for this effect, we regress individual




Figure A.4: Rising unemployment duration: distribution and alternative explanations

Source: Current Population Survey, Employment Opportunity Pilot Project (1979-80), Survey of Consumer Expectations (2013-22), and [Landais, Michailat, and Saez \(2018\)](#). *Notes:* The upper left panel plots the shares unemployed for less than one month and more than six months (left axis), alongside mean elapsed duration (right axis), excluding temporary layoffs. The upper right panel compares raw and demographically adjusted mean duration, normalized to the same initial level. The lower left panel plots the effective UI replacement rate. The lower right panel compares densities of demographically-adjusted log hourly reservation wages in 2019 dollars across the EOPP and SCE; vertical lines mark the means.

elapsed unemployment duration on education and demographics such as age, gender and race and plot the residualized unemployment duration in the upper right panel of Figure A.4. We normalize the residualized series to match the level of raw unemployment duration at the start of the sample. Between the expansions of the late 1990s and the late 2010s, the raw average rises from 3.0 to 4.8 months and the residualized average from 2.6 to 3.9 months. Composition shifts thus account for roughly a quarter of the increase,

and the remainder reflects longer spells within demographic groups.

Another potential explanation is that more generous unemployment insurance (UI) has prolonged job search. The lower-left panel reproduces the effective UI replacement rate, capturing the level and duration of benefits, constructed by [Landaï, Michailat, and Saez \(2018\)](#) for 1990-2014. The series fluctuates countercyclically around a mean of 42% but shows no sustained upward trend, ruling out increasing UI generosity as a potential driver for rising unemployment duration over the years shown.

A third potential explanation is that workers have become more selective about the jobs they accept and are willing to wait longer in search of a better offer. We examine this using the EOPP and SCE surveys, both of which ask job seekers about their reservation wages. The EOPP asks respondents who are not working and looking for work, *“What is the lowest hourly wage rate you would accept?”*. The SCE asks, *“Suppose someone offered you a job today. What is the lowest wage or salary you would accept, before taxes and deductions, for the type of work you are looking for?”*, and we keep respondents who were unemployed in the month before the interview. We convert answers into hourly rates and express reservation wages in 2019 dollars, take logs, and trim the top and bottom percentile within each survey. We residualize log real reservation wages with respect to age, squared age, gender and three education groups, and add back the survey-specific weighted mean.

The lower right panel plots the kernel densities across the two surveys (EOPP 1979-80 and SCE 2013-22). The two distributions overlap closely, and there is little change in their means indicated by the dashed lines. The rise in unemployment duration therefore does not appear to reflect a general increase in worker selectivity. However, the distribution in 2013-22 is notably less concentrated and has a wider tails, indicating a widening of dispersion in reservation wages.⁴⁰


A.5 Diverging duration: construction and robustness

In this appendix, we describe how Figure 5 is constructed and show that the divergence in unemployment duration across worker types is robust to comparing workers at the same age rather than the same calendar year.

We use the cross-sectional samples of the NLSY79, which follows a cohort born in 1957 to 1964, and of the NLSY97, which follows a cohort born in 1980 to 1984. We drop the supplemental and military samples and keep the weeks in which a respondent is

⁴⁰This widening of the distribution of reservation is in line with a mechanism rooted in increasing worker heterogeneity, such as the one explored in our model.

(a) By lifetime average unemployment duration



(b) By AFQT score




Figure A.5: Mean unemployment duration by worker age

Source: National Longitudinal Survey of Youth 1979 and 1997. *Notes:* Each panel plots the mean duration of completed unemployment spells among workers averaged over four year age bins. Red lines refer to the NLSY79 cohort and blue lines to the NLSY97 cohort. In the left panel, dashed lines (“above median”) refer to workers whose lifetime average unemployment duration is above the cohort median, solid lines (“below median”) refer to those at or below it. In the right panel, dashed lines (“low AFQT”) refer to workers in the bottom AFQT quintile and solid lines (“high AFQT”) to workers in the top AFQT quintile.

between 25 and 64 years old. The samples contain 6,086 and 6,682 respondents, of whom 3,947 and 3,430 experience at least one completed unemployment spell. They contribute 13,082 and 7,676 spells, completed between 1982 and 2022 in the NLSY79 and between 2005 and 2021 in the NLSY97. Both surveys record each worker’s labor force status for every week since the first interview.

We define an unemployment spell as a period in a worker’s labor market history that begins with a week of unemployment, continues through any subsequent weeks of non-participation, and ends when the respondent is employed for at least two consecutive weeks. We compute its duration as the number of weeks of unemployment within the spell. We keep only completed spells whose beginning and transition to employment are observed and drop spells of a single week. Each spell is assigned to the year in which it ends.

In any NLSY cohort, as calendar time and age advance together, the lengthening of low-type spells could in principle reflect a life-cycle component of duration: rather than capturing a divergence of spell length over time, Figure 5 may thus be capturing that




Figure A.6: Online job search and labor market outcomes across states

Source: Current Population Survey, Job Openings and Labor Turnover Survey, and Lightcast. *Notes:* Each circle represents a state. Online exposure is the log deviation of a state's share of Lightcast postings to its share of JOLTS vacancies in 2019. The panels report log changes from 2000 to 2023. Circle sizes are proportional to initial employment, and gray lines are employment-weighted linear fits.

spell durations naturally diverge when cohorts age. To distinguish age-effects from time-effects, the left and right panels of Figure A.5 thus plot the classifications by lifetime average unemployment duration and AFQT scores against age. The left panel shows that workers with below median lifetime average unemployment duration have spells of about two months at every age in both cohorts. In the NLSY79, spells of workers above the median lengthen over the life cycle from about five months in the late twenties to thirteen months in the mid-fifties. In contrast, in the NLSY97, workers above the median already reach ten to twelve months in their early thirties, roughly twice the duration of the NLSY79 workers at the same age. The cross-cohort comparison thus reveals that time rather than age effects play a central role in the divergence of unemployment spell lengths.

A.6 Additional cross-sectional evidence

Cross-state validation. Figure A.6 repeats our cross-sectoral analysis from Section 1.5 at the state level. We construct exposure to online vacancies as each state's share of Lightcast online job postings relative to its share of JOLTS job openings, using average monthly counts in 2019. We compute the state-level job filling rate as the ratio of hires to vacancies from the JOLTS data, and the job finding rate as the monthly transition probability from unemployment to employment, as well as average unemployment duration from the CPS. We use the same comparison periods as in the sectoral analysis from the main

text, circle sizes are the initial state-level employment and fitted lines are correspondingly weighted regression lines. The figure shows that states more exposed to online job search experienced larger declines in job filling and finding rates and larger increases in unemployment duration. The patterns across sectors therefore also appear across states.

German Job Vacancy Survey We use the fourth quarter waves of the IAB Job Vacancy Survey (JVS) for 2005-21. The JVS draws a representative cross-section of German establishments with at least one employee subject to social security contributions. Our analysis uses the module on the establishments most recent hire into employment, and describes their entire recruitment process.

Our main independent variable of interest is whether the establishment searched online for its most recent hire. We classify a job search as online if the establishment advertised through its own website, an online job board, the Federal Employment Agency's websites, or social media. Offline methods include newspaper advertisements, private placement agencies, personal contacts, unsolicited applications, and non-internet contact with the Federal Employment Agency. Establishments that combined online and offline methods belong to the online group. Searches reported only through internal recruitment or an unclassified other channel do not enter the comparison.

We construct five outcomes. Applications are the total number of applicants reported for the position, which we take in logs conditional on a positive count. Days to fill measures the time between the beginning of search and the hiring decision. We drop inconsistent dates and searches lasting more than two years. We capture screening effort through two measures. First, we compute the total working hours spent filling the position. These hours cover the entire recruitment process, from preparing the advertisement through evaluating candidates, interviewing, and finalizing the employment contract. The second measure counts assessment methods used beyond interviews and application materials, including tests, trial work, references, internet searches, medical examinations, criminal-record checks, etc. An explicit response that no additional method was used is coded as zero. Finally, the long-term unemployment indicator equals one if the hired worker had been unemployed for more than one year and zero if unemployment lasted up to one year. This outcome is defined only among hires from unemployment.

The outcomes are available in different waves. Applications and days to fill cover 2005-21. Recruitment hours are recorded in 2014-15 and 2017-21. Assessment methods are recorded in 2018. Prior unemployment duration is available from 2008 onward. These differences account for the different sample sizes across columns.

	Screening effort				
	Log applications (1)	Log days to fill (2)	Log screening hours (3)	Assessment methods (4)	Pr(hiring LTU) (5)
Searched online	0.927*** (0.021)	0.499*** (0.021)	0.722*** (0.028)	0.279*** (0.049)	−0.055** (0.020)
Year FE	Y	Y	Y	Y	Y
Observations	93,325	87,637	43,621	5,384	19,660
R^2	0.169	0.067	0.109	0.028	0.017

Table 5: Online search and recruitment outcomes in German establishments

Source: IAB Job Vacancy Survey, 2005-2021. *Notes:* Each column reports a separate regression of the dependent variable (denoted by the column name) on an indicator for whether the establishment used at least one online search channel for its most recent hire. LTU denotes long-term unemployed. All specifications include year fixed effects. Outcomes are available in different survey waves (applications and days to fill 2005-21; screening hours 2014-15 and 2017-21, assessment methods 2018; and LTU hiring 2008-21) accounting for the differences in sample size. Standard errors, clustered by establishment, in parentheses. Establishments weights used. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

All regressions use establishment survey weights and standard errors clustered by establishment. The specification in Table 1 of the main text include year fixed effects, required qualifications, broad time-consistent occupation and industry categories, log employment, and the log gross hiring rate. The latter is the number of hires over the preceding twelve months divided by current employment. Required qualifications are grouped into three levels: unskilled or short training, vocational or technical qualifications, and college or university qualifications. Table 5 reports estimates with year fixed effects but without vacancy and establishment controls, using the same observations as the main text specifications. All five coefficients retain their signs.

A.7 Additional aggregate-series validation

Duration and inverse job-finding hazards. Figure A.7 provides an alternative way of examining the heterogeneity in job finding rates among the unemployed documented in Figure 5. The left panel plots two measures of the monthly probability of transitioning out of unemployment following Shimer (2012): a flows-based Unemployment-to-Employment (UE) transition probability (dashed green) and a stock-based outflow prob-




Figure A.7: Unemployment duration and unemployment exit rates

Source: Current Population Survey and authors' calculations. *Notes:* The left panel compares the flow-based probability of moving from unemployment to employment (dashed green) with the stock-based probability of exiting unemployment (solid blue). The right panel compares mean elapsed unemployment duration (dashed red) with the inverse of the outflow rate (solid blue), normalized to match duration at the start of the sample. All series are smoothed using centered twelve-month moving averages.

ability (blue).⁴¹ By construction, the stock-based outflow probability is broader, capturing all exits from unemployment, to either employment or non-participation, so it exceeds the UE probability. The figure shows that both series begin exhibiting a decline starting in the early 2000s. The outflow probability averaged 44% between 1998-1999 and decreased to 34% in the 2018-19 period. The UE probability displays a similar albeit less pronounced decline, trending down from an average of nearly 28% in the late 1990s to 27% in the late-2010s.

The right panel compares the mean unemployment duration with the inverse of the outflow rate, the latter being normalized to the level of the former at the beginning of the sample period. The comparison can be interpreted through the Jensen's inequality. In

⁴¹The UE probability is the fraction of unemployed in month $t - 1$ who transition to employment in month t , UE_t/U_{t-1} . The outflow probability F_t is derived from the law of motion of unemployment:

$$U_{t+1} = (1 - F_t)U_t + U_{t+1}^{<5w}$$

where U_t is the unemployment stock and $U_{t+1}^{<5w}$ approximates new entrants to unemployment using the stock of unemployed less than five weeks in $t + 1$.




Figure A.8: Matching efficiency and online job search

Source: Current Population Survey, Job Openings and Labor Turnover Survey, Lightcast, and [Petrosky-Nadeau and Zhang \(2021\)](#). *Notes:* Matching efficiency is defined as the hires from unemployment divided by the square root of the product of unemployment and vacancies. The left panel plots the log of matching efficiency as a four-quarter moving average, normalized to zero at the start of the sample. The middle and right panels relate log changes in matching efficiency across sectors and states to online exposure, using the comparison periods, exposure definitions, and employment weights in Figures 6 and A.6.

steady state, with constant worker-specific job finding rate f_i , average elapsed duration is $\mathbb{E}[1/f_i]$ and the inverse of the aggregate job finding rate is $1/\mathbb{E}[f_i]$, where both expectations are taken over the unemployed. Since $1/f$ is convex, $\mathbb{E}[1/f_i] \geq 1/\mathbb{E}[f_i]$, with equality only if all unemployed workers face the same job finding rate. The gap between average duration and the inverse job finding rate is therefore indicative of the dispersion of job finding rates across the unemployed. The figure shows that the gap between the two series stays stable until the 2001 recession, and then widens substantially during the Great Recession, remaining elevated through the subsequent expansion. The widening gap thus reflects growing dispersion in job finding prospects that Figure 5 documents in the NLSY data.

Speaking to the decline of matching efficiency. While we focus on the matching rates of firms and workers separately in the main text, we now show that our empirical results can alternatively be framed in terms of matching efficiency, a common measure used to assess changes in the rate at which the labor market produces matches between job seekers and open positions. We define matching efficiency in its simplest form, assuming a Cobb-Douglas meeting technology that takes as inputs vacancies and unemployed job seekers, and gives as output UE hires:

$$ME_u = \frac{UE \text{ hires}}{\sqrt{\text{Vacancies} \cdot \text{Unemployed}}} \quad (7)$$

This definition is often motivated as the residual of a matching function: how many hires do we see in the data relative the number of hires we should expect for a given number of unemployed job seekers and vacancies?

The left panel of Figure A.8 plots matching efficiency, as defined by equation (7), which we report in logs and normalize to zero at the start of the sample. During the early decades, the series remains roughly stable around zero, showing no discernible trend. Beginning in the late 1990s, a gradual downward trend starts setting in, and after the Great Recession matching efficiency starts declining precipitously. The pandemic sets in motion a brief rebound but by the end of the sample matching efficiency settles at a level about 0.5 log points below its level in 1975. This implies that, for a given stock of job seekers and vacancies, the labor market now produces about 40% fewer hires than it did five decades ago. The middle and right panels of Figure A.8 document that the excess decline in matching efficiency is more pronounced in sectors and states that are relatively more exposed to online job search.

Controlling for the business cycle. We show that the increase in average unemployment duration, and the decline in job finding and filling rates, and matching efficiency were not a result of slack labor markets or the aggregate state of the business cycle. To do this, we plot all aggregate series in the seven expansionary episodes of the last 50 years against labor market tightness (Figure A.9) and unemployment rate (Figure A.10) on the horizontal axis. Figure A.9 panels (a) and (b) show that unemployment duration shifted upward and job finding probability shifted downward relative to their prior recovery paths after the Great Recession, even at high levels of labor market tightness. The job filling rate (panel c) and matching efficiency (panel d) shifted down over time. Figure A.10 shows similar broad patterns using the unemployment rate as an indicator of the business cycle.

B Model appendix

B.1 Alternative formulation of the screening technology

In the main text, we assume that a firm draws a cost of screening a worker and may then decide to screen or not conditional on the realization of the cost draw. Here, we




Figure A.9: Labor market outcomes and tightness across recoveries

Source: Current Population Survey, Job Openings and Labor Turnover Survey, [Petrosky-Nadeau and Zhang \(2021\)](#). *Notes:* Each panel plots the indicated variable against the vacancy-to-unemployment ratio, with both axes in logs. Colors identify recovery episodes. Duration and UE probability use centered twelve-month moving averages, and matching efficiency uses a four-quarter moving average. Unemployment duration is shown as a six-month lead.

demonstrate that this model is equivalent to one where the firm decides to optimally set a probability of screening subject to a convex cost.

We begin by setting up the alternative firm problem. Let Δ^f denote the return to screening, as it does in the main text. In our alternative formulation, we assume that the cost of




Figure A.10: Labor market outcomes and unemployment across recoveries

Source: Current Population Survey, Job Openings and Labor Turnover Survey, [Petrosky-Nadeau and Zhang \(2021\)](#). Notes: Each panel plots the indicated variable against the unemployment rate, with both axes in logs. Colors identify recovery episodes. Duration and UE probability use centered twelve-month moving averages, and matching efficiency uses a four-quarter moving average. Unemployment duration is shown as a six-month lead.

screening with probability p^{sig} is

$$\phi(p^{\text{sig}}) = \gamma_s^{-1} \frac{\eta_s}{1 + \eta_s} (p^{\text{sig}})^{\frac{1+\eta_s}{\eta_s}}$$

and thus, defining $\tilde{q} = q(a_L^u p_L^{\text{app}} + (1 - a_L^u) p_H^{\text{app}})$ as the rate at which firms encounter applicants, the value function of an unmatched vacancy is

$$(r + \delta)V = \max_{p^{\text{sig}}} \tilde{q}(p^{\text{sig}} \Delta^f - \phi(p^{\text{sig}}))$$

and the first order condition of this problem is

$$\begin{aligned} \Delta^f &= \phi'(p^{\text{sig}}) = \gamma_s^{-1} (p^{\text{sig}})^{\frac{1}{\eta_s}} \\ \implies p^{\text{sig}} &= \gamma_s (\Delta^f)^{\eta_s} = G(\Delta^f) \end{aligned}$$

which is the same relationship between the return to screening and the probability of screening that is featured in the baseline model. Thus, the two models are equivalent.

B.2 Value functions and equilibrium definition

In this appendix, we complete the formal description of all decision problems in our model and define an equilibrium.

B.2.1 Surplus, value functions and optimal decisions

Surplus function and wage setting. The surplus function for a match of type i is defined as

$$S_i = W_i - U_i + J_i, \tag{8}$$

where U_i and W_i , respectively, denote the value of unemployment and employment to a type- i worker. J_i is the value to the firm of forming a match with type- i worker.⁴²

We assume that wages are set by Nash bargaining – that is, they solve

$$w_i = \arg \max_{w_i} J_i(w_i)^{1-\beta} (W_i(w_i) - U_i)^\beta.$$

Thus, the surplus accruing to the firm and worker in a match is, respectively,

$$W_i - U_i = \beta S_i, \quad J_i = (1 - \beta) S_i. \tag{9}$$

⁴²Note that the vacancy value does not enter the surplus expression because we assume that filled jobs cannot costlessly return to search. The consequence of this assumption is that it is never optimal for the firm to fire a type- L worker and return to search.

We now describe the value functions of workers and firms and derive optimal application and screening decisions.

Value of an unemployed worker. The value to an unemployed worker of type i , U_i , satisfies

$$rU_i = b_i + \lambda \mathbb{E} \left[\max \left\{ 0, \underbrace{\left[p^{\text{sig}} \left(\mathbb{I}_{i=L}(1-s) + \mathbb{I}_{i=H}s \right) + (1-p^{\text{sig}})(1-R) \right]}_{p_i^{\text{acc}} \equiv \text{Acceptance probability}} (W_i - U_i) - c \right\} \right] \quad (10)$$

An unemployed worker receives a flow value of b_i and meets vacancies at rate λ . Upon meeting, the worker draws an application cost c and decides whether to apply. Equation (10) captures this decision through the max operator: the worker applies whenever the expected gain from an application, which we denote as $\Delta_i^w \equiv p_i^{\text{acc}} \cdot (W_i - U_i)$, exceeds the cost c . The endogenous probability of applying is therefore given by $p_i^{\text{app}} = F(\Delta_i^w)$. The acceptance probability p_i^{acc} is the sum of two terms. If the firm obtains a signal about the worker, which occurs with probability p^{sig} , then type- L workers are accepted with probability $1-s$ and type- H workers with probability s .⁴³ If the firm does not obtain a signal, which occurs with probability $1-p^{\text{sig}}$, the firm's hiring decision does not depend on the worker's true type. We define $R \in [0, 1]$ as the probability that the firm rejects applicants in this scenario.

In what follows, a useful object is the worker's job finding rate $f_i = \lambda F(\Delta_i^w) p_i^{\text{acc}}$. It is the product of the meeting rate, the application probability and the acceptance probability.

Value of an employed worker. The value of an employed worker of type i , is characterized by the following Bellman equation:

$$rW_i = w_i + \delta(U_i - W_i). \quad (11)$$

An employed worker receives a flow wage w_i and faces exogenous separation at rate δ , upon which the worker returns to unemployment.

⁴³This follows from an argument we outline below: the optimal acceptance strategy of a firm that has purchased a signal is to admit the worker if and only if the signal is H .

Value of a filled job. A firm matched to a worker of type i receives a value J_i where

$$(r + \delta)J_i = y_i - w_i \quad (12)$$

The flow value of a filled job is the output y_i produced by the worker net of the wage w_i . The match dissolves at rate δ , at which point the firm exits with zero value.

Value of an unfilled job. The value of an unfilled job opening, V , satisfies the following Bellman equation:

$$(r + \delta)V = \underbrace{q(a_L^u p_L^{\text{app}} + (1 - a_L^u) p_H^{\text{app}})}_{\text{encounter rate}} \cdot \mathbb{E} \left[\max\{V_{\text{signal}} - V - \kappa, V_{\text{nosignal}} - V\} \right]. \quad (13)$$

An unfilled job meets applicants at rate q . The unemployed pool is composed of type- L workers (with share a_L^u) and type- H workers who apply with probabilities p_L^{app} and p_H^{app} , respectively. Conditional on receiving an application, the firm draws a screening cost κ and decides whether to purchase a signal. The max operator captures this decision: the firm screens whenever the net gain from obtaining the signal, $V_{\text{signal}} - V_{\text{nosignal}}$, exceeds the cost κ . We define this threshold as $\Delta^f = V_{\text{signal}} - V_{\text{nosignal}}$, so the probability of purchasing a signal is $p^{\text{sig}} = G(\Delta^f)$.

The firm's hiring decision depends on whether it has purchased a signal. Consider first the case in which it has. As we show in Appendix B.3, a firm that purchases a signal optimally accepts the worker if and only if the signal is H . The expected value to a firm that has purchased the signal is therefore

$$V_{\text{signal}} = \underbrace{(1 - a_L)sJ_H + a_L(1 - s)J_L}_{\text{average value of accepting}} + \underbrace{\left(1 - (1 - a_L)s - a_L(1 - s)\right)V}_{\text{average value of rejecting}}, \quad (14)$$

where a_L is the share of type- L workers among all applicants that we characterize in Section B.2.2. The first term captures the case in which the worker is truly type H and the signal is correct: this occurs with probability $(1 - a_L)s$ and yields value J_H . The second term captures the case in which the worker is truly type L but the signal is incorrect: this occurs with probability $a_L(1 - s)$ and yields value J_L . In the remaining cases, the signal is L and the firm rejects the applicant, retaining the vacancy value V .

We now consider the case in which the firm decides not to purchase a signal. We

express the value of a firm as:

$$V_{\text{nosignal}} = \max_R (1 - R) \cdot [(1 - a_L)J_H + a_L J_L] + R \cdot V = \max\{V, (1 - a_L)J_H + a_L J_L\}. \quad (15)$$

Without a signal, the firm cannot distinguish between worker types and faces a binary choice. It can accept the applicant, receiving the expected match value, or reject the applicant and remain vacant. By accepting applicants without a signal, the firm fills the vacancy more quickly but with a worker of lower expected productivity. By rejecting such applicants, the firm preserves the option to fill the job with an applicant of higher expected productivity later on. The rejection probability without a signal, R , captures how the firm resolves this trade-off. When the expected value of an unscreened hire, $(1 - a_L)J_H + a_L J_L$, exceeds the vacancy value V , the firm accepts all unscreened applicants ($R = 0$). When it falls below, the firm rejects them ($R = 1$). For R to lie in the interior, the firm must be indifferent between accepting and rejecting workers without a signal.

Free entry. To close the firm side of the model, we turn to entry. To create a new vacancy, a potential entrant must pay a sunk entry cost k . Free entry implies that entrants create vacancies (and thereby change q) until the value of a vacancy is driven to the entry cost, so that $V = k$ in equilibrium.

B.2.2 Flows

Having characterized the optimal decisions of workers and firms, we now derive the steady-state conditions that pin down the aggregate composition of the unemployed and the applicant pool. For type- i workers the outflow from unemployment must equal inflow:

$$u_L f_L = \delta(\alpha_L^{\text{pop}} - u_L), \quad u_H f_H = \delta((1 - \alpha_L^{\text{pop}}) - u_H) \quad (16)$$

where u_i denotes the steady-state mass of unemployed type- i workers.

The share of type- L workers among the unemployed follows directly from the steady-state unemployment levels:

$$a_L^u \equiv \frac{u_L}{u_L + u_H} = \frac{\alpha_L^{\text{pop}}(f_H + \delta)}{\alpha_L^{\text{pop}} f_H + (1 - \alpha_L^{\text{pop}})f_L + \delta} \quad (17)$$

Similarly, the share of type- L workers among applicants depends on the composition of the unemployed pool and the type-specific application probabilities:

$$a_L = \frac{a_L^u p_L^{\text{app}}}{a_L^u p_L^{\text{app}} + (1 - a_L^u) p_H^{\text{app}}} \quad (18)$$

B.2.3 Equilibrium

Definition (equilibrium).

An *equilibrium* is a tuple $(\lambda, q, V, V_{\text{signal}}, V_{\text{nosignal}}, \{W_i, U_i, J_i, S_i\}_{i \in \{L, H\}}, \{w_i\}_{i \in \{L, H\}}, R, a_L^u, a_L)$, such that the above Bellman equations ((8), (9), (10), (11), (12), (13), (14), (15)) and flow equations ((16), (17), (18)) are satisfied, and λ and q are consistent with the meeting function (equation (2)).

As discussed above, the firm's rejection decision without a signal, R , governs the key trade-off faced by the firm: filling a job quickly but with workers of lower expected productivity, or filling it slowly but with workers of higher expected productivity. Since changes in this trade-off will be central to the mechanism we study, we focus on equilibria in which firms choose $R \in (0, 1)$. That is, the firm follows a mixed strategy when deciding whether to reject applicants without a signal.

Definition (mixed-strategy equilibrium).

A *mixed-strategy equilibrium* is an equilibrium that features $R \in (0, 1)$.

B.3 Optimal hiring following a signal

In this appendix, we show that any firm purchasing a signal hires if and only if it receives an H -signal. We assume that a firm does not purchase a signal if doing so has no value, so a firm that purchases a signal pays $\kappa > 0$.

Posterior beliefs. Let a_L^{app} denote the share of type- L workers among applicants. Bayes' rule gives

$$\begin{aligned}\Pr(H \mid H\text{-signal}) &= \frac{(1 - a_L^{\text{app}})s}{(1 - a_L^{\text{app}})s + a_L^{\text{app}}(1 - s)}, \\ \Pr(H \mid \text{no signal}) &= 1 - a_L^{\text{app}}, \\ \Pr(H \mid L\text{-signal}) &= \frac{(1 - a_L^{\text{app}})(1 - s)}{(1 - a_L^{\text{app}})(1 - s) + a_L^{\text{app}}s}.\end{aligned}$$

Since $s \geq 0.5$, these probabilities satisfy

$$\Pr(H \mid H\text{-signal}) \geq \Pr(H \mid \text{no signal}) \geq \Pr(H \mid L\text{-signal}). \quad (19)$$

Ordering of match values. We next show that $J_H \geq J_L > 0$ in any equilibrium. As we show in Appendix B.4, the value equations and Nash bargaining imply

$$J_i = \frac{(1 - \beta)\rho y_i}{r + \delta + \beta(1 + \eta_a)^{-1}f_i} > 0. \quad (20)$$

Suppose, by contradiction, that $J_H < J_L$. Since $J_i = (1 - \beta)S_i$, this implies $S_H < S_L$.

Let x^H , x^L , and x^0 denote the firm's hiring probabilities following an H -signal, an L -signal, and no signal, respectively. Under $J_H < J_L$, the expected value of hiring is decreasing in the posterior probability that the applicant is type H . Equation (19) therefore implies that the expected value of hiring is weakly higher following an L -signal than following an H -signal. Because the firm compares both expected hiring values with the same vacancy value V , optimality requires $x^L \geq x^H$.

The type-specific acceptance probabilities are

$$\begin{aligned}p_L^{\text{acc}} &= (1 - p^{\text{sig}})x^0 + p^{\text{sig}}[sx^L + (1 - s)x^H], \\ p_H^{\text{acc}} &= (1 - p^{\text{sig}})x^0 + p^{\text{sig}}[(1 - s)x^L + sx^H].\end{aligned}$$

Consequently,

$$p_L^{\text{acc}} - p_H^{\text{acc}} = p^{\text{sig}}(2s - 1)(x^L - x^H) \geq 0. \quad (21)$$

Using $W_i - U_i = \beta S_i$, equations (20) and (21), together with $S_H < S_L$, imply

$$p_H^{\text{acc}}(W_H - U_H) = \beta p_H^{\text{acc}} S_H \leq \beta p_L^{\text{acc}} S_L = p_L^{\text{acc}}(W_L - U_L). \quad (22)$$

Because F is increasing, equation (22) implies

$$F(p_H^{\text{acc}}(W_H - U_H)) \leq F(p_L^{\text{acc}}(W_L - U_L)). \quad (23)$$

It follows that

$$f_H = \lambda p_H^{\text{acc}} F(p_H^{\text{acc}}(W_H - U_H)) \leq \lambda p_L^{\text{acc}} F(p_L^{\text{acc}}(W_L - U_L)) = f_L. \quad (24)$$

Using equation (20), $y_H \geq y_L$, and $f_H \leq f_L$ then gives

$$\frac{J_H}{J_L} = \frac{y_H}{y_L} \frac{r + \delta + \beta(1 + \eta_a)^{-1} f_L}{r + \delta + \beta(1 + \eta_a)^{-1} f_H} \geq 1, \quad (25)$$

contradicting $J_H < J_L$. Therefore, $J_H \geq J_L$ in every equilibrium.

Hiring following an H -signal. For each information state $\sigma \in \{H, 0, L\}$, define the expected value of hiring as

$$\bar{J}_\sigma = \Pr(H | \sigma) J_H + [1 - \Pr(H | \sigma)] J_L,$$

where $\sigma = 0$ denotes the absence of a signal. The ordering of match values and equation (19) imply

$$\bar{J}_H \geq \bar{J}_0 \geq \bar{J}_L. \quad (26)$$

Consider a firm that purchases a signal at cost $\kappa > 0$. If $\bar{J}_H \leq V$, equation (26) implies $\bar{J}_0 \leq V$ and $\bar{J}_L \leq V$. The firm would then retain the vacancy following either signal realization and would also retain it without purchasing a signal, contradicting its willingness to pay $\kappa > 0$. Hence,

$$\bar{J}_H > V,$$

so any firm that purchases a signal hires after receiving an H -signal.

Rejection following an L -signal. Suppose that a firm purchasing a signal at cost $\kappa > 0$ optimally hires after receiving an L -signal. Optimality then requires $\bar{J}_L \geq V$. Equation (26) implies

$$\bar{J}_H \geq \bar{J}_0 \geq \bar{J}_L \geq V.$$

The firm would therefore hire following either signal realization and would also hire without purchasing a signal. Moreover, the law of iterated expectations gives

$$\Pr(H\text{-signal})\bar{J}_H + \Pr(L\text{-signal})\bar{J}_L = \bar{J}_0.$$

Thus, purchasing the signal and hiring following either realization yields exactly the same expected gross value as hiring without purchasing the signal. The signal would have zero gross value, contradicting the firm's willingness to pay $\kappa > 0$. Hence,

$$\bar{J}_L < V,$$

so any firm that purchases a signal rejects after receiving an L -signal. Therefore, a firm that purchases a signal hires if and only if the signal is H .

B.4 Derivation of mixed-strategy equilibrium conditions

In this appendix, we define the functions referenced by the screening curve (SC) and indifference curve (IC) from the main text, as well as a number of auxiliary functions that are useful for those definitions. All functions we define are functions of $\vec{f} = (f_L, f_H)$ that take parameters as given, and are defined recursively. By this we mean that, in what follows below, every equation encodes instructions on how to compute a new object using only $\vec{f}$, parameters, and previously defined objects. Thus, all objects defined in equations (27)-(35) are closed-form functions of $\vec{f}$. One implication of this is that, fixing some parameters, knowing the solution $\vec{f}$ of the system defined by equations (IC) and (SC) is enough to derive all equilibrium objects defined below in closed form. The functions are defined as follows:

$$J_i(f_i) = \left(\frac{(1-\beta)\rho y_i}{r + \delta + \beta(1+\eta_a)^{-1}f_i} \right) \quad (27)$$

$$a_L^u(\vec{f}) = \frac{\alpha_L^{\text{pop}}(f_H + \delta)}{\alpha_L^{\text{pop}}f_H + (1 - \alpha_L^{\text{pop}})f_L + \delta} \quad (28)$$

$$a_L(\vec{f}) = a_L^u(\vec{f}) \left(a_L^u(\vec{f}) + (1 - a_L^u(\vec{f})) \left(\frac{f_H}{f_L} \right)^{\frac{\eta_a}{1+\eta_a}} \left(\frac{r + \delta + \beta(1+\eta_a)^{-1}f_L}{r + \delta + \beta(1+\eta_a)^{-1}f_H} \right)^{\frac{\eta_a}{1+\eta_a}} \left(\frac{y_H}{y_L} \right)^{\frac{\eta_a}{1+\eta_a}} \right)^{-1} \quad (29)$$

$$\Delta^f(\vec{f}) = a_L(\vec{f})(1 - a_L(\vec{f}))(1 - \beta)(2s - 1).$$

$$\left(\frac{1}{r + \delta + \beta(1+\eta_a)^{-1}f_H}(\rho y_H) - \frac{1}{r + \delta + \beta(1+\eta_a)^{-1}f_L}(\rho y_L) \right) \quad (30)$$

$$\lambda(\vec{f}) = \left(\frac{Z}{((r + \delta)k)\omega} \right)^{\frac{(1+\eta_a)}{1+\eta_a-\omega}} \left[\gamma_a^{\frac{1}{1+\eta_a}} (1 + \eta_a)^{\frac{\eta_a}{1+\eta_a}} \left(a_L^u(\vec{f}) \left(\frac{\beta(1+\eta_a)^{-1}f_L}{r + \delta + \beta(1+\eta_a)^{-1}f_L}(\rho y_L) \right)^{\frac{\eta_a}{1+\eta_a}} \right. \right. \\ \left. \left. + (1 - a_L^u(\vec{f})) \left(\frac{\beta(1+\eta_a)^{-1}f_H}{r + \delta + \beta(1+\eta_a)^{-1}f_H}(\rho y_H) \right)^{\frac{\eta_a}{1+\eta_a}} \right) (1 + \eta_s)^{-1} \gamma_s \Delta^f(\vec{f})^{(1+\eta_s)} \right]^{\frac{\omega(1+\eta_a)}{1+\eta_a-\omega}} \quad (31)$$

$$\Delta_L^w(\vec{f}) = \left(\gamma_a^{-1}(1 + \eta_a)\lambda(\vec{f})^{-1} \left(\frac{\beta(1+\eta_a)^{-1}f_L}{r + \delta + \beta(1+\eta_a)^{-1}f_L} \right) (\rho y_L) \right)^{\frac{1}{1+\eta_a}} \quad (32)$$

$$\Delta_H^w(\vec{f}) = \left(\gamma_a^{-1}(1 + \eta_a)\lambda(\vec{f})^{-1} \left(\frac{\beta(1+\eta_a)^{-1}f_H}{r + \delta + \beta(1+\eta_a)^{-1}f_H} \right) (\rho y_H) \right)^{\frac{1}{1+\eta_a}} \quad (33)$$

$$p_i^{\text{app}}(\vec{f}) = F(\Delta_i^w(\vec{f})) \quad (34)$$

$$p^{\text{sig}}(\vec{f}) = G(\Delta^f(\vec{f})) \quad (35)$$

In what follows, we derive these equations alongside the screening and indifference curves from the main text.

Worker side. With $F(c) = \gamma_a c^{\eta_a}$,

$$\int_0^\Delta c dF(c) = \gamma_a \eta_a \int_0^\Delta c^{\eta_a} dc = \frac{\eta_a}{1 + \eta_a} F(\Delta) \Delta,$$

and hence

$$\begin{aligned} rU_i &= b + \lambda \int_0^{\Delta_i^w} (\Delta_i^w - c) dF(c) = b + \lambda \left(\Delta_i^w F(\Delta_i^w) - \int_0^{\Delta_i^w} c dF(c) \right) \\ &= b + \lambda F(\Delta_i^w) (1 + \eta_a)^{-1} \Delta_i^w. \end{aligned}$$

Plugging this into the surplus Bellman equation, we get

$$(r + \delta)S_i = y_i - rU_i = y_i - b - \frac{\lambda\gamma_a}{1 + \eta_a}(\Delta_i^w)^{1+\eta_a}.$$

We know that by Nash bargaining and the definition of Δ_i^w , $S_i = \Delta_i^w / (\beta p_i^{\text{acc}})$. Substituting this into the previous equation, we get

$$(r + \delta) \frac{\Delta_i^w}{\beta p_i^{\text{acc}}} = y_i - b - \frac{\lambda\gamma_a}{1 + \eta_a}(\Delta_i^w)^{1+\eta_a}.$$

Since $f_i = \lambda F(\Delta_i^w) p_i^{\text{acc}} = \lambda \gamma_a (\Delta_i^w)^{\eta_a} p_i^{\text{acc}}$, we can eliminate p_i^{acc} and write

$$(r + \delta) \Delta_i^w = \frac{\beta f_i}{\lambda \gamma_a (\Delta_i^w)^{\eta_a}} \rho y_i - \frac{\beta f_i}{1 + \eta_a} \Delta_i^w.$$

Solving for Δ_i^w yields

$$\Delta_i^w = \left[\gamma_a^{-1} (1 + \eta_a) \lambda^{-1} \frac{\beta (1 + \eta_a)^{-1} f_i}{r + \delta + \beta (1 + \eta_a)^{-1} f_i} \rho y_i \right]^{\frac{1}{1+\eta_a}}.$$

which establishes equations (32), (33) and (34). We can immediately use this result to derive

$$\frac{F(\Delta_H^w)}{F(\Delta_L^w)} = \left(\frac{f_H}{f_L} \right)^{\frac{\eta_a}{1+\eta_a}} \left(\frac{r + \delta + \beta (1 + \eta_a)^{-1} f_L}{r + \delta + \beta (1 + \eta_a)^{-1} f_H} \right)^{\frac{\eta_a}{1+\eta_a}} \left(\frac{y_H}{y_L} \right)^{\frac{\eta_a}{1+\eta_a}}. \quad (36)$$

which will be useful below.

Firm side. From $J_i = (1 - \beta)S_i$ and the surplus equation one obtains the compact form

$$J_i = \frac{(1 - \beta) \rho y_i}{r + \delta + \beta (1 + \eta_a)^{-1} f_i}.$$

which establishes equation (27). We now derive Δ^f . The value when screening an applicant is

$$V_{\text{signal}} = (1 - a_L)sJ_H + a_L(1 - s)J_L + \left(1 - (1 - a_L)s - a_L(1 - s) \right) V.$$

For a mixed-strategy equilibrium, the firm is indifferent between rejecting and accepting an unscreened worker, which implies

$$V_{\text{nosignal}} = (1 - a_L)J_H + a_L J_L = V = k.$$

Compute the return to screening from its definition:

$$\begin{aligned}\Delta^f &= V_{\text{signal}} - V_{\text{nosignal}} \\ &= (1 - a_L)s(J_H - V) + a_L(1 - s)(J_L - V).\end{aligned}$$

Write $J_H - V$ and $J_L - V$ as

$$\begin{aligned}J_H - V &= J_H - ((1 - a_L)J_H + a_L J_L) = a_L(J_H - J_L), \\ J_L - V &= J_L - ((1 - a_L)J_H + a_L J_L) = -(1 - a_L)(J_H - J_L).\end{aligned}$$

Substitute into Δ^f :

$$\begin{aligned}\Delta^f &= (1 - a_L)s \cdot a_L(J_H - J_L) + a_L(1 - s) \cdot (-(1 - a_L)(J_H - J_L)) \\ &= a_L(1 - a_L)(J_H - J_L)(s - (1 - s)) \\ &= a_L(1 - a_L)(2s - 1)(J_H - J_L).\end{aligned}$$

Using the Bellman equations for J_L and J_H yields

$$\Delta^f = a_L(1 - a_L)(1 - \beta)(2s - 1) \left[\frac{\rho y_H}{r + \delta + \beta(1 + \eta_a)^{-1}f_H} - \frac{\rho y_L}{r + \delta + \beta(1 + \eta_a)^{-1}f_L} \right].$$

which establishes equations (30) and (35).

The aggregate encounter rate of the worker. Free entry implies

$$(r + \delta)k = q \left(a_L F(\Delta_L^w) + (1 - a_L) F(\Delta_H^w) \right) \frac{\gamma_s}{1 + \eta_s} (\Delta^f)^{1 + \eta_s}.$$

Solving for q :

$$q = \frac{(r + \delta)k}{\left(a_L F(\Delta_L^w) + (1 - a_L)F(\Delta_H^w)\right)(1 + \eta_s)^{-1} \gamma_s(\Delta^f)^{1+\eta_s}}.$$

Now use $\lambda = Z^{\frac{1}{1-\omega}} q^{\frac{\omega}{\omega-1}}$ and plug in Δ_i^w from above:

$$\lambda^{1+\frac{\omega\eta_a}{(1+\eta_a)(1-\omega)}} = \left(\frac{Z}{((r + \delta)k)^\omega}\right)^{\frac{1}{1-\omega}} \left[\gamma_a^{\frac{1}{1+\eta_a}} (1 + \eta_a)^{\frac{\eta_a}{1+\eta_a}} \left(a_L^u A_L + (1 - a_L^u) A_H\right) (1 + \eta_s)^{-1} \gamma_s(\Delta^f)^{1+\eta_s} \right]^{\frac{\omega}{1-\omega}}.$$

where

$$A_i = \left(\frac{\beta(1 + \eta_a)^{-1} f_i}{r + \delta + \beta(1 + \eta_a)^{-1} f_i} \rho y_i \right)^{\frac{\eta_a}{1+\eta_a}}.$$

Solving this equation for λ gives

$$\lambda = \left(\frac{Z}{((r + \delta)k)^\omega} \right)^{\frac{1+\eta_a}{1+\eta_a-\omega}} \left[\gamma_a^{\frac{1}{1+\eta_a}} (1 + \eta_a)^{\frac{\eta_a}{1+\eta_a}} \left(a_L^u A_L + (1 - a_L^u) A_H\right) (1 + \eta_s)^{-1} \gamma_s(\Delta^f)^{1+\eta_s} \right]^{\frac{\omega(1+\eta_a)}{1+\eta_a-\omega}}.$$

which establishes equation (31).

Flows and stocks. We now derive the flow identities. Let u_i be the steady-state mass unemployed of type i and $u \equiv u_L + u_H$. Stationarity implies

$$u_i f_i = \delta(\alpha_i^{\text{pop}} - u_i) \quad \Rightarrow \quad u_i = \frac{\delta \alpha_i^{\text{pop}}}{f_i + \delta}.$$

Hence the share of L among the unemployed is

$$a_L^u = \frac{u_L}{u_L + u_H} = \frac{\alpha_L^{\text{pop}}(f_H + \delta)}{\alpha_L^{\text{pop}} f_H + (1 - \alpha_L^{\text{pop}}) f_L + \delta},$$

which establishes equation (28). The share of L -types among applicants follows Bayes' Law:

$$a_L = \frac{a_L^u F(\Delta_L^w)}{a_L^u F(\Delta_L^w) + (1 - a_L^u) F(\Delta_H^w)} = \frac{a_L^u}{a_L^u + (1 - a_L^u) \frac{F(\Delta_H^w)}{F(\Delta_L^w)}},$$

which, together with equation (36), establishes equation (29).

Screening curve. Start from the definition of type-specific job finding rates

$$f_i = \lambda F(\Delta_i^w) p_i^{\text{acc}} \quad (i \in \{L, H\}).$$

The acceptance probabilities satisfy

$$\begin{aligned} p_L^{\text{acc}} &= (1 - G(\Delta^f))(1 - R) + G(\Delta^f)(1 - s) \\ p_H^{\text{acc}} &= (1 - G(\Delta^f))(1 - R) + G(\Delta^f)s \\ \implies p_H^{\text{acc}} - p_L^{\text{acc}} &= G(\Delta^f)(s - (1 - s)) = G(\Delta^f)(2s - 1). \end{aligned}$$

Now express f_H in terms of f_L :

$$\begin{aligned} f_H &= \lambda p_H^{\text{app}} p_H^{\text{acc}} \\ &= \lambda p_H^{\text{app}} \left(p_L^{\text{acc}} + (p_H^{\text{acc}} - p_L^{\text{acc}}) \right) \\ &= \lambda p_H^{\text{app}} p_L^{\text{acc}} + \lambda p_H^{\text{app}} p^{\text{sig}}(2s - 1). \end{aligned}$$

Replace $p_L^{\text{acc}} = f_L / (\lambda F(\Delta_L^w))$ (from the definition of f_L) to obtain

$$f_H = \frac{p_H^{\text{app}}}{p_L^{\text{app}}} f_L + \lambda p_H^{\text{app}} p^{\text{sig}}(2s - 1).$$

This is the screening curve (SC) used in the main text.

Indifference curve (IC). By equation (15), in a mixed-strategy equilibrium, the value of accepting an unscreened applicant equals the vacancy value V , which by free entry equals k in equilibrium:

$$V = k = (1 - a_L)J_H + a_L J_L.$$

This is the indifference curve (IC) used in the main text.

B.5 Slope of the Indifference Curve

Let J'_i denote the derivative of J_i with respect to f_i . Taking the total derivative along the indifference curve gives

$$0 = \left[a_L J'_L + (J_L - J_H) \frac{\partial a_L}{\partial f_L} \right] df_L + \left[(1 - a_L) J'_H + (J_L - J_H) \frac{\partial a_L}{\partial f_H} \right] df_H.$$

Using $J_H > J_L$, $J'_i < 0$, $\partial a_L / \partial f_L < 0$, and $\partial a_L / \partial f_H > 0$, the slope of the indifference curve can be written as

$$\left. \frac{df_H}{df_L} \right|_{IC} = \frac{(J_H - J_L) \left(-\frac{\partial a_L}{\partial f_L} \right) - a_L (-J'_L)}{(1 - a_L) (-J'_H) + (J_H - J_L) \frac{\partial a_L}{\partial f_H}}.$$

Every term in the denominator is positive. The sign of the slope is therefore determined by the difference between the two positive terms in the numerator. Two forces determine the slope: first, a higher f_L decreases a_L , so f_H must rise to keep the firm indifferent; second, a higher f_L decreases J_L , so f_H must fall to keep the firm indifferent. On the left-hand side of the IC, a_L is close to one and changes little with f_L , so the second term dominates and the indifference curve slopes downward. Further to the right, a_L responds more strongly to f_L , strengthening the first effect, which eventually induces a moderate upward slope.

C Calibration appendix

This appendix provides further context and details on the calibration choices and targets. We begin by discussing why Z and s can be normalized in the model and then provide some background for the empirical series used to anchor our targets.

C.1 Normalizations

We begin by discussing why Z and s can be normalized.⁴⁴

⁴⁴The case of y_L is obvious and we omit it here; it represents a unit normalization that is common to search models.

Normalization of Z . Z is a normalization because Z and γ_a are not separately identified. To see this, let $\theta = v/u$ denote market tightness, so that $\lambda = Z\theta^\omega$ and $q = Z\theta^{\omega-1}$. Consider an alternative parametrization with $Z' = \tau Z$ and $\gamma'_a = \gamma_a/\tau$ for some $\tau > 0$. We guess and verify that market tightness θ remains unchanged. At unchanged θ , this implies $\lambda' = \tau\lambda$, $q' = \tau q$, and $F'(c) = F(c)/\tau$. Letting $\bar{c}_i$ denote the type- i application threshold, it follows that

$$\lambda' \int_0^{\bar{c}_i} (\bar{c}_i - c) dF'(c) = \lambda \int_0^{\bar{c}_i} (\bar{c}_i - c) dF(c), \quad q' p_i^{\text{app}'} = q p_i^{\text{app}}.$$

Thus, workers' value equations and application thresholds remain unchanged, while firms face the same applicant-arrival rates and applicant composition. The same values, policies, job-finding rates, job-filling rates, and stocks therefore constitute an equilibrium and market tightness remains the same.

The model is thus observationally equivalent to one in which encounters occur more frequently but are less likely to produce an application. Consequently, Z can be normalized to any value consistent with $p_i^{\text{app}} < 1$. Choosing a high value ensures that this condition holds.

Normalization of s . As we show in Proposition 2, the equilibrium conditions determining observable outcomes depend on s and γ_s only through the composite

$$\gamma_s(2s - 1)^{1+\eta_s}.$$

Replacing s by any $s' \in (1/2, 1]$ and setting

$$\gamma'_s = \gamma_s \left(\frac{2s - 1}{2s' - 1} \right)^{1+\eta_s}$$

leaves the equilibrium allocation unchanged. The screening probability changes, but the no-signal rejection probability R adjusts so that type-specific acceptance rates remain unchanged. Thus, s and γ_s are not separately identified. Intuitively, the model is observationally equivalent to one in which signals are less costly but also less informative.

The choice of s does affect the implied value of R , which must remain between zero and one to constitute a mixed-strategy equilibrium. We set $s = 1$ to keep the parameter space for which mixed-strategy equilibria exist maximally large. Any normalization of s sufficiently close to one delivers an observationally equivalent model with a mixed

strategy equilibrium.

C.2 Empirical targets

Here, we provide additional background information on the empirical targets we choose to target.

Coefficient of variation of wages. We choose the coefficient of variation of wages, defined as the ratio between the standard deviation of wages and their mean, to match the first row of Table 1 in [Hornstein et al. \(2011\)](#), which corresponds to a value of 0.47.

Job finding rate, job filling rate, and unemployment duration. We use the series from Section 1 to obtain targets for the job finding rate, job filling rate, and unemployment duration of the late 1990s. Concretely, we take averages from December 1998 to December 2000. The reason for this choice of time frame is that it reflects a two-year period during which the business cycle was at a peak, with low unemployment rates. The variables therefore offer a useful benchmark when we compare them in the data and the model to the post-internet time period of December 2017 to December 2019, which likewise reflected a business cycle peak.

Measuring average applications in the EOPP and SCE. We measure average applications submitted in the last four weeks by unemployed job seekers aged 21-65. In the EOPP (1979-80), we first sum job applications across search methods within each search spell. We then prorate each spell's applications by its overlap with the last 28 days before the interview. Finally, we average the prorated applications among the unemployed respondents. In the SCE (2013-19, 2021-22) we average jobs applied to in the last four weeks among respondents who were unemployed in the previous month. Job seekers with zero applications are included in both averages. Please refer to Appendix A.1 for more information about each of the surveys.

Measuring the relative elasticity of job finding rates. We use η_a to target differences in the cyclicalities of job finding rates across worker types. To get an empirical estimate of this moment, we classify NLSY79 respondents into quintiles based on their average unemployment-spell duration over the sample, assigning workers in the longest-duration

quintile to type L and workers in the remaining four quintiles to type H .⁴⁵ For each quintile q and year t , we compute the mean weekly unemployment-to-employment transition probability, f_{qt}^w , and convert it to a monthly probability as

$$F_{qt} = 1 - (1 - f_{qt}^w)^{52/12}.$$

We then estimate separately by quintile

$$\log F_{qt} = \alpha_q + \varepsilon_q \log u_t + \beta_{1q}t + \beta_{2q}t^2 + e_{qt},$$

where u_t is the annual average unemployment rate from the BLS. We define the elasticity for type H as the average coefficient across the four shorter-duration quintiles, $\hat{\varepsilon}_H = \frac{1}{4} \sum_{q=1}^4 \hat{\varepsilon}_q$, and the elasticity for type L as the coefficient for the longest-duration quintile. Our target is their difference, $\hat{\varepsilon}_H - \hat{\varepsilon}_L = 0.763$.⁴⁶ In the model, we construct the corresponding moment by increasing productivity by 1% and recomputing the steady-state equilibrium. Let $f_i = 1 - \exp(-\lambda_i)$ denote the monthly job finding probability of type i , let u denote the aggregate steady-state unemployment rate, and let primes denote the equilibrium after the productivity shock. The model counterpart is

$$\frac{[\log f'_H - \log f_H] - [\log f'_L - \log f_L]}{\log u' - \log u},$$

which measures the difference in the unemployment elasticities of job finding rates between the two worker types.

Measuring the cost of screening in the NES. We construct the screening resource per hire from the last-hired-worker supplement of the 1982 National Employer Survey, which asks about the position of the last employee hired before August 1981. The survey reports total personnel hours spent on recruiting, screening, and interviewing all applicants for that position, and the number of applicants who accepted the job. Our sample contains 2,044 establishments and we compute screening hours per hire as the ratio of survey-weighted personnel hours spent to survey-weighted total hires, obtaining 8.31 hours per hire. Assuming a 40-hour workweek, monthly working hours equal $40 \times 52/12 = 173.3$, so the total screening resources per hire correspond to 4.8% of a

⁴⁵This is broadly consistent with the model, which has a share of L -type workers of approximately 24%.

⁴⁶This is in the same range as the results of a similar exercise in [Mann \(2026\)](#).

full-time worker's monthly hours.

C.3 Mapping Cowgill et al. (2024) into the model

Cowgill et al. (2024) run a screening experiment in which “senders” produce short written pitches and “receivers” evaluate them. Some senders have access to an LLM when writing the pitch, while receivers do not. Receivers then report a posterior belief $p \in [0, 1]$ that the sender is an “expert” (where true expertise is a binary variable). The authors report that $p(1 - p)$ increases from 0.17 (without LLMs) to 0.18 (with LLMs), indicating that signals contained in the pitch are less informative when the LLM is available.

To map this to our model, we interpret the firm's signaling technology as analogous to the “pitch” technology in CHW's experiment. Given a symmetric prior, the firm's posterior probability of the worker being H upon receiving an H signal is s , which implies that p and s can be treated as analogous. The reported change of $p(1 - p)$ implies that

$$\frac{2p' - 1}{2p - 1} = \frac{2\sqrt{0.25 - p'(1 - p')}}{2\sqrt{0.25 - p(1 - p)}}$$

and therefore the corresponding proportional change in $2s - 1$ is

$$\frac{2s^{\text{AI}} - 1}{2s^{\text{noAI}} - 1} = \frac{2\sqrt{0.25 - 0.18}}{2\sqrt{0.25 - 0.17}} = \sqrt{\frac{0.07}{0.08}} \approx 0.935,$$

which corresponds to a 6.5% decline in $2s - 1$.

D Derivations of model formulae and proofs

D.1 Encounter rate

For the equilibria we focus on, the encounter rate is pinned down by the indifference condition of the firm

$$k = V = (r + \delta)^{-1} \tilde{q} G(\Delta^f) (\Delta^f - \mathbb{E}[\kappa \mid \kappa \leq \Delta^f])$$

where $\Delta^f = V_{\text{signal}} - V_{\text{nosignal}} = a_L(1 - a_L)(2s - 1)(J_H - J_L)$ and

$$\tilde{q} = q(a_L^u p_L^{\text{app}} + (1 - a_L^u) p_H^{\text{app}})$$

is the rate at which firms encounter applicants. Given the functional form $G(x) = \gamma_s x^{\eta_s}$,

$$G(\Delta^f)(\Delta^f - \mathbb{E}[\kappa \mid \kappa \leq \Delta^f]) = (1 + \eta_s)^{-1} G(\Delta^f) \Delta^f,$$

and thus,

$$\begin{aligned} k(r + \delta) &= (1 + \eta_s)^{-1} \tilde{q} G(\Delta^f) \Delta^f \\ &= (1 + \eta_s)^{-1} \tilde{q} \gamma_s^{-\frac{1}{\eta_s}} (p^{\text{sig}})^{\frac{1+\eta_s}{\eta_s}}, \end{aligned} \quad (37)$$

which implies

$$\tilde{q} = k(r + \delta) (1 + \eta_s) \gamma_s^{1/\eta_s} (p^{\text{sig}})^{-\frac{1+\eta_s}{\eta_s}}. \quad (38)$$

Equation (38) illustrates how the encounter rate $\tilde{q}$ is determined in equilibrium: Since only encounters during which a signal is obtained produce value in excess of k to the firm, the encounter rate of firms is solely determined by the excess value of such events, Δ^f . This has parallels to the free-entry logic of a standard DMP model: a higher value of encounters implies lower encounter rates for firms.

D.2 Unemployment duration decomposition

In the model, the average duration of unemployment is

$$\text{unemployment duration} = a_L^u \frac{1}{f_L} + (1 - a_L^u) \frac{1}{f_H}$$

and the aggregate job-finding rate is

$$\text{job finding rate} = (a_L^u f_L + (1 - a_L^u) f_H).$$

Therefore,

$$\begin{aligned}
\frac{\text{unemployment duration}}{(\text{job finding rate})^{-1}} &= \left(a_L^u \frac{1}{f_L} + (1 - a_L^u) \frac{1}{f_H} \right) \cdot \left(a_L^u f_L + (1 - a_L^u) f_H \right) \\
&= \left((a_L^u)^2 + (1 - a_L^u)^2 + a_L^u (1 - a_L^u) \left(\frac{f_H}{f_L} + \frac{f_L}{f_H} \right) \right) \\
&= \left(1 + a_L^u (1 - a_L^u) \left(\frac{f_H}{f_L} + \frac{f_L}{f_H} - 2 \right) \right) \\
&= \left(1 + a_L^u (1 - a_L^u) \frac{(f_H - f_L)^2}{f_H f_L} \right). \tag{39}
\end{aligned}$$

D.3 Pre-separation wage rank of hires

To construct the pre-separation wage rank, we first calculate each worker type's rank in the steady-state wage distribution. Since the model has two wage levels, the average rank of L -type and H -type workers is $\alpha_L^e/2$ and $(1 + \alpha_L^e)/2$, respectively, where α_L^e is the type- L share of employment.

For each type $i \in \{L, H\}$, we then evolve an initially employed population for 12 months according to the type-specific transition matrix:

$$\dot{x}_{it} = M_i x_{it} \implies x_{it} = \exp(M_i t) x_{i0}$$

where

$$M_i = \begin{bmatrix} -\delta & f_i \\ \delta & -f_i \end{bmatrix}$$

and the state is $x_{it} = (e_{it}, u_{it})'$ with $x_{i0} = (\alpha_i^e, 0)$.

Let $u_{i,12}$ denote the resulting mass unemployed after 12 months, then $u_{i,12} f_i$ is proportional to the current hiring flow among workers with a wage record 12 months earlier. Letting $\pi_{L,12} = \frac{u_{L,12} f_L}{u_{L,12} f_L + u_{H,12} f_H}$ denote the type- L share of this hiring flow, the model-implied average pre-separation wage rank is

$$\mathcal{R}^{\text{wage}} = \pi_L^{12} \frac{\alpha_L^e}{2} + \left(1 - \pi_L^{12} \right) \frac{1 + \alpha_L^e}{2}.$$

D.4 Unemployment duration rank of hires

Let $a_L^{\text{hire}} \in [0, 1]$ denote the share of type- L workers among hires. The unemployment-duration distribution in the unemployed stock is

$$F_{\text{dur}}(x) = 1 - a_L^u e^{-f_L x} - (1 - a_L^u) e^{-f_H x},$$

and the unemployment-duration density among hires is

$$f_{\text{hire}}(x) = a_L^{\text{hire}} f_L e^{-f_L x} + (1 - a_L^{\text{hire}}) f_H e^{-f_H x}.$$

The average percentile rank (CDF-rank) of a typical hire's unemployment duration within the unemployed distribution is

$$\mathcal{R} := \mathbb{E}[F_{\text{dur}}(X_{\text{hire}})] = \int_0^\infty F_{\text{dur}}(x) f_{\text{hire}}(x) dx.$$

Expanding and integrating term-by-term,

$$\begin{aligned} \mathcal{R} &= \int_0^\infty \left(a_L^{\text{hire}} f_L e^{-f_L x} + (1 - a_L^{\text{hire}}) f_H e^{-f_H x} \right) dx \\ &\quad - a_L^u \int_0^\infty e^{-f_L x} \left(a_L^{\text{hire}} f_L e^{-f_L x} + (1 - a_L^{\text{hire}}) f_H e^{-f_H x} \right) dx \\ &\quad - (1 - a_L^u) \int_0^\infty e^{-f_H x} \left(a_L^{\text{hire}} f_L e^{-f_L x} + (1 - a_L^{\text{hire}}) f_H e^{-f_H x} \right) dx \\ &= 1 - a_L^u \left(\frac{a_L^{\text{hire}}}{2} + \frac{(1 - a_L^{\text{hire}}) f_H}{f_L + f_H} \right) - (1 - a_L^u) \left(\frac{a_L^{\text{hire}} f_L}{f_L + f_H} + \frac{1 - a_L^{\text{hire}}}{2} \right) \\ &= 1 - \frac{a_L^u a_L^{\text{hire}}}{2} - \frac{(1 - a_L^u)(1 - a_L^{\text{hire}})}{2} - \frac{a_L^u (1 - a_L^{\text{hire}}) f_H + (1 - a_L^u) a_L^{\text{hire}} f_L}{f_L + f_H}. \end{aligned}$$

The share of type- L workers among hires satisfies

$$a_L^{\text{hire}} = \frac{a_L^u f_L}{a_L^u f_L + (1 - a_L^u) f_H}, \quad 1 - a_L^{\text{hire}} = \frac{(1 - a_L^u) f_H}{a_L^u f_L + (1 - a_L^u) f_H}.$$

Therefore,

$$\frac{a_L^u a_L^{\text{hire}}}{2} = \frac{(a_L^u)^2 f_L}{2(a_L^u f_L + (1 - a_L^u) f_H)},$$

$$\frac{(1 - a_L^u)(1 - a_L^{\text{hire}})}{2} = \frac{(1 - a_L^u)^2 f_H}{2(a_L^u f_L + (1 - a_L^u) f_H)},$$

and

$$a_L^u(1 - a_L^{\text{hire}})f_H + (1 - a_L^u)a_L^{\text{hire}}f_L = \frac{a_L^u(1 - a_L^u)(f_H^2 + f_L^2)}{a_L^u f_L + (1 - a_L^u) f_H}.$$

Substituting these expressions gives

$$\mathcal{R} = 1 - \frac{(a_L^u)^2 f_L + (1 - a_L^u)^2 f_H}{2(a_L^u f_L + (1 - a_L^u) f_H)} - \frac{a_L^u(1 - a_L^u)(f_H^2 + f_L^2)}{(f_L + f_H)(a_L^u f_L + (1 - a_L^u) f_H)}.$$

Using the identity

$$(a_L^u)^2 f_L + (1 - a_L^u)^2 f_H = a_L^u f_L + (1 - a_L^u) f_H - a_L^u(1 - a_L^u)(f_L + f_H),$$

we obtain

$$\mathcal{R} = \frac{1}{2} + \frac{a_L^u(1 - a_L^u)}{a_L^u f_L + (1 - a_L^u) f_H} \left[\frac{f_L + f_H}{2} - \frac{f_H^2 + f_L^2}{f_L + f_H} \right].$$

Finally,

$$\frac{f_L + f_H}{2} - \frac{f_H^2 + f_L^2}{f_L + f_H} = -\frac{(f_H - f_L)^2}{2(f_L + f_H)},$$

which yields

$$\mathcal{R} = \frac{1}{2} \left(1 - \frac{a_L^u(1 - a_L^u)(f_H - f_L)^2}{(f_L + f_H)(a_L^u f_L + (1 - a_L^u) f_H)} \right).$$

A few observations on $\mathcal{R}$ are in order. First, without heterogeneity, $\mathcal{R} = \frac{1}{2}$. Second, $\mathcal{R} \leq \frac{1}{2}$ for any (f_L, f_H) , indicating that the average duration rank of hires is no greater than the average duration rank in the unemployed pool, which by definition equals $\frac{1}{2}$. Lastly, as job-finding rates become more dispersed, the average duration rank of hires falls.

D.5 Rate of screening expenditures

In this appendix, we show that the rate of screening resource consumption depends only on k , r , δ , and η_s . We first note that the rate of screening is defined by

$$\text{screening expenditure rate} = \tilde{q} \cdot \mathbb{E} \left[\kappa \cdot \mathbb{I}\{\kappa \leq \Delta^f\} \right]$$

where $\tilde{q} = q(a_L^u p_L^{\text{app}} + (1 - a_L^u) p_H^{\text{app}})$ is the rate at which vacancies encounter applicants. Substituting in the free entry condition from equation (37) in Appendix D.1 yields

$$\text{screening expenditure rate} = k(r + \delta)\eta_s. \quad (40)$$

Therefore, the rate at which firms spend resources on screening does not depend on changes in equilibrium that are driven by parameters other than k , r , δ , and η_s .

D.6 Proofs of propositions

Proof of Proposition 1. We start by defining

$$\hat{\Delta}^f(\vec{f}) := a_L(\vec{f}) \left(1 - a_L(\vec{f})\right) (1 - \beta) \left[\frac{\rho y_H}{r + \delta + \beta(1 + \eta_a)^{-1} f_H} - \frac{\rho y_L}{r + \delta + \beta(1 + \eta_a)^{-1} f_L} \right],$$

which yields

$$\begin{aligned} \Delta^f(\vec{f}) &= (2s - 1) \hat{\Delta}^f(\vec{f}) \\ \implies G(\Delta^f(\vec{f}))(2s - 1) &= \gamma_s (2s - 1)^{1 + \eta_s} \left(\hat{\Delta}^f(\vec{f}) \right)^{\eta_s}. \end{aligned} \quad (41)$$

Next, define

$$\hat{\lambda}(\vec{f}) := \left(\frac{Z}{((r + \delta)k)^\omega} \right)^{\frac{1 + \eta_a}{1 + \eta_a - \omega}} \left[(1 + \eta_a)^{\frac{\eta_a}{1 + \eta_a}} \Gamma(\vec{f}) (1 + \eta_s)^{-1} \left(\hat{\Delta}^f(\vec{f}) \right)^{1 + \eta_s} \right]^{\frac{\omega(1 + \eta_a)}{1 + \eta_a - \omega}},$$

where

$$\begin{aligned} \Gamma(\vec{f}) &:= a_L^u(\vec{f}) \left[\frac{\beta(1 + \eta_a)^{-1} f_L}{r + \delta + \beta(1 + \eta_a)^{-1} f_L} (\rho y_L) \right]^{\frac{\eta_a}{1 + \eta_a}} \\ &\quad + \left(1 - a_L^u(\vec{f})\right) \left[\frac{\beta(1 + \eta_a)^{-1} f_H}{r + \delta + \beta(1 + \eta_a)^{-1} f_H} (\rho y_H) \right]^{\frac{\eta_a}{1 + \eta_a}}, \end{aligned}$$

and write

$$\begin{aligned}\lambda(\vec{f}) &= \left(\frac{Z}{((r+\delta)k)^\omega} \right)^{\frac{1+\eta_a}{1+\eta_a-\omega}} \left[\gamma_a^{\frac{1}{1+\eta_a}} (1+\eta_a)^{\frac{\eta_a}{1+\eta_a}} \Gamma(\vec{f}) (1+\eta_s)^{-1} \gamma_s \left((2s-1) \hat{\Delta}^f(\vec{f}) \right)^{1+\eta_s} \right]^{\frac{\omega(1+\eta_a)}{1+\eta_a-\omega}} \\ &= \hat{\lambda}(\vec{f}) \gamma_a^{\frac{\omega}{1+\eta_a-\omega}} \gamma_s^{\frac{\omega(1+\eta_a)}{1+\eta_a-\omega}} (2s-1)^{\frac{\omega(1+\eta_a)(1+\eta_s)}{1+\eta_a-\omega}}.\end{aligned}\quad (42)$$

Combining the second term of the screening curve, we get

$$\begin{aligned}& \lambda(\vec{f}) F(\Delta_w^H(\vec{f})) G(\Delta^f(\vec{f})) (2s-1) \\ &= \lambda(\vec{f}) \cdot \left[\gamma_a^{\frac{1}{1+\eta_a}} (1+\eta_a)^{\frac{\eta_a}{1+\eta_a}} \lambda(\vec{f})^{-\frac{\eta_a}{1+\eta_a}} \left(\frac{\beta(1+\eta_a)^{-1} f_H}{r+\delta+\beta(1+\eta_a)^{-1} f_H} (\rho y_H) \right)^{\frac{\eta_a}{1+\eta_a}} \right] \\ & \quad \cdot \left[\gamma_s (2s-1)^{1+\eta_s} (\hat{\Delta}^f(\vec{f}))^{\eta_s} \right] \\ &= \gamma_a^{\frac{1}{1+\eta_a}} \gamma_s (2s-1)^{1+\eta_s} \lambda(\vec{f})^{\frac{1}{1+\eta_a}} \cdot (1+\eta_a)^{\frac{\eta_a}{1+\eta_a}} \left(\frac{\beta(1+\eta_a)^{-1} f_H}{r+\delta+\beta(1+\eta_a)^{-1} f_H} (\rho y_H) \right)^{\frac{\eta_a}{1+\eta_a}} (\hat{\Delta}^f(\vec{f}))^{\eta_s}.\end{aligned}$$

Use equation (42) to substitute in for $\lambda(\vec{f})$:

$$\lambda(\vec{f})^{\frac{1}{1+\eta_a}} = \hat{\lambda}(\vec{f})^{\frac{1}{1+\eta_a}} \gamma_a^{\frac{\omega}{(1+\eta_a)(1+\eta_a-\omega)}} \gamma_s^{\frac{\omega}{1+\eta_a-\omega}} (2s-1)^{\frac{\omega(1+\eta_s)}{1+\eta_a-\omega}}.$$

This finally yields

$$\lambda(\vec{f}) F(\Delta_w^H(\vec{f})) G(\Delta^f(\vec{f})) (2s-1) = X(\gamma_a, \gamma_s, s)^{\frac{1}{1+\eta_a-\omega}} \cdot M_2(\vec{f}),$$

where

$$X(\gamma_a, \gamma_s, s) = \gamma_a^{1+\eta_a} \gamma_s^{1+\eta_s} (2s-1)^{(1+\eta_s)(1+\eta_a)}.$$

and

$$M_2(\vec{f}) := (1+\eta_a)^{\frac{\eta_a}{1+\eta_a}} \left(\frac{\beta(1+\eta_a)^{-1} f_H}{r+\delta+\beta(1+\eta_a)^{-1} f_H} (\rho y_H) \right)^{\frac{\eta_a}{1+\eta_a}} (\hat{\Delta}^f(\vec{f}))^{\eta_s} \hat{\lambda}(\vec{f})^{\frac{1}{1+\eta_a}},$$

which does not depend on (γ_a, γ_s, s) .

Finally, $M_1(\vec{f}) = \frac{F(\Delta_H^w)}{F(\Delta_L^w)}$, which does not depend on (γ_a, γ_s, s) by equation (36). This

establishes the claim. $\square$

Proof of Proposition 2. Starting from the result of Proposition 1, if $\vec{f}'$ constitutes an equilibrium, then the indifference curve holds and

$$f'_H = M_1(\vec{f}') f'_L + X(\gamma'_a, \gamma_s, s)^{\frac{1}{1+\eta_a-\omega}} \cdot M_2(\vec{f}').$$

Note that neither $M_1(\vec{f})$ nor the indifference curve depend on (γ_a, γ_s, s) . Thus, if

$$\gamma'_s = \gamma_s \left(\frac{\gamma'_a}{\gamma_a} \right)^{\frac{1}{1+\eta_a}}, \quad (43)$$

$$s' = \frac{1}{2} \left[1 + (2s - 1) \left(\frac{\gamma'_a}{\gamma_a} \right)^{\frac{1}{(1+\eta_s)(1+\eta_a)}} \right], \quad (44)$$

then the indifference curve holds and

$$\begin{aligned} f'_H &= M_1(\vec{f}') f'_L + X(\gamma_a, \gamma'_s, s)^{\frac{1}{1+\eta_a-\omega}} \cdot M_2(\vec{f}') \\ &= M_1(\vec{f}') f'_L + X(\gamma_a, \gamma_s, s')^{\frac{1}{1+\eta_a-\omega}} \cdot M_2(\vec{f}') \end{aligned}$$

and thus the screening curve holds.

We still need to verify that there is ε such that $R \in (0, 1)$ for any $(f'_L, f'_H) \in B_\varepsilon(\vec{f})$. This claim follows directly from the fact that (f_L, f_H) are part of a mixed-strategy equilibrium and the fact that the expression

$$R(x) = 1 - \frac{\frac{f_L}{\lambda(x) F(\Delta_w^L(x))} - G(\Delta^f(x)) (1 - s)}{1 - G(\Delta^f(x))}.$$

is continuous in $x = (\vec{f}, \gamma_a, \gamma_s, s)$. Therefore, the screening curve and indifference condition hold at $\vec{f}'$ for E' , E'_s , and E'_{γ_s} and $R \in (0, 1)$ holds true in each case. This establishes the claim that the equilibria exist and feature the same type-specific job finding rates. From this claim it follows directly that the distribution of unemployment duration is identical across all economies.

Next, we establish that all economies also feature the same job filling rates, and the same unemployment and vacancy stocks. As established above, $\vec{f}' = \vec{f}'_s = \vec{f}'_{\gamma_s}$, and hence the type-specific unemployment stocks (and total unemployment) are identical across the

three equilibria. Moreover, combining (42) with (43) and (44) implies that $\lambda(\vec{f}')$ is identical across E' , E'_s , and E'_{γ_s} . From $\lambda = Z \cdot (v/u)^\omega$, it follows that tightness v/u is identical, and therefore the vacancy stock v is identical as well. Total hires are $H = \sum_{i \in \{L, H\}} u_i f_i$, which is identical since both u_i and f_i are identical. Thus the job filling rate H/v is identical across the three equilibria. The same applies to match efficiency, since H , v and u are all identical.

We now show that E'_s and E'_{γ_s} feature the same type-specific application rates ($\lambda F(\Delta_w^L)$, $\lambda F(\Delta_w^H)$). Given the insight that λ is constant across E' , E'_{γ_s} and E'_s , this claim follows directly from the fact that the expression for $F(\Delta_w^H)$ depends on s and γ_s only through λ .

Finally, we prove the claim that the screening rates are identical across E'_{γ_a} and E'_{γ_s} . This can be shown by observing that

$$k(r + \delta) = \tilde{q}(1 + \eta_s)^{-1} G(\Delta^f) \Delta^f$$

which implies that

$$k(r + \delta)(1 + \eta_s)(\Delta^f)^{-1} = \underbrace{\tilde{q} \cdot G(\Delta^f)}_{\text{Screening rate}} .$$

The claim then follows from the observation that, by equation (41), Δ^f depends on s , but not on γ_a and γ_s . □