AER: Insights 2021, 3(3): 339–352
https://doi.org/10.1257/aeri.20200576

Jacks of All Trades and Masters of One: Declining Search Frictions and Unequal Growth†

By Paolo Martellini and Guido Menzio*

Declining search frictions generate productivity growth by allowing workers to find jobs for which they are better suited. For "jacks of all trades"—workers whose productivity is similar across different jobs in their labor market—declining search frictions lead to minimal growth. For "masters of one trade"—workers whose productivity varies a great deal across different jobs in their labor market—declining search frictions lead to fast growth. A rudimentary calibration suggests that differential returns to declining search frictions may account for a non-negligible fraction of the wage growth differential between routine and nonroutine workers. (*JEL* J24, J31, J63, J64, O33)

Search frictions cause labor misallocation in the sense that workers are not necessarily employed in the jobs where they are most productive. As search frictions decline—due to, say, advances in information and communication technologies—the allocation of labor improves and average labor productivity increases.[1] The return of declining search frictions on productivity, though, is unequal across different groups of workers. For workers who are generalists or "jacks of all trades," in the sense that their productivity is similar across different jobs in their labor market, the decline in search frictions leads to minimal productivity and wage growth. For workers who are specialists or "masters of one trade"—in the sense that their productivity varies a great deal across different jobs—the decline in search frictions leads to high productivity and wage growth.

We formalize the above argument using a version of the search-theoretic model by Martellini and Menzio (2020, henceforth MM) with multiple types of workers. Different types of workers populate different labor markets. In each market, workers

* Martellini: Federal Reserve Bank of Minneapolis (email: pa.martellini@gmail.com); Menzio: Department of Economics, New York University (email: gm1310@nyu.edu). Pete Klenow was the coeditor for this article. We thank the three referees for their suggestions. We are grateful to Martin Beraja, Pieter Gautier, Victoria Gregory, Ilse Lindenlaub, Simon Mongey, Chris Moser, Robert Ulbricht, and Gianluca Violante for their help. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

† Go to https://doi.org/10.1257/aeri.20200576 to visit the article page for additional materials and author disclosure statement(s).

[1] Bhuller, Kostøl, and Vigtel (2019) provide direct evidence that information technology lowers search frictions. They show that, after the introduction of broadband internet, firms find employees more easily, workers face fewer problems locating available jobs, and starting wages are higher. A lumpy fall of search frictions in our model generates precisely these effects.

and firms are located along a circle of unit circumference and meet through a random search process. When a worker and a firm meet, their distance along the circumference (which we interpret as the distance between the skills of the worker and the skill requirements of the firm) determines their match-specific component of productivity and, in turn, their decision of whether to consummate the match. The environment is nonstationary. Specifically, the efficiency of the production technology and the efficiency of the search process grow over time at some constant, exogenous rate.

As in MM, we restrict attention to a balanced growth path (BGP), an equilibrium consistent with the empirical observation that the unemployment rate, the vacancy rate, the rate at which employed workers become unemployed (EU rate), and the rate at which unemployed workers become employed (UE rate) have no secular trend. The existence of a BGP implies some restrictions on the fundamentals of the model. In particular, it implies that the match-specific component of productivity must be an isoelastic function of the firm-worker skill distance.

Our main theoretical finding is to show that, for each type of worker, there are two components to the growth rate of productivity and wages. The first component is the growth rate of the efficiency of the production technology. The second component is the growth rate of the efficiency of the search process multiplied by the elasticity of the match-specific productivity with respect to the firm-worker skill distance. Therefore, the productivity growth for two types of workers may differ not only because of biased technical change, i.e., the efficiency of the production technology used by one type of worker grows faster than the efficiency of the production technology used by the other type. The productivity growth for two types of workers may differ also because the return to declining search frictions is higher for one type—the type with a higher elasticity of the match-specific productivity with respect to skill distance (i.e., the masters of one trade)—and lower for the other type—the type with a lower elasticity of the match-specific productivity with respect to skill distance (i.e., the jacks of all trades). We then show that the return to declining search frictions for a particular type of worker is related to the extent of cross-sectional productivity (and, in turn, wage) dispersion in that type's labor market. Thus, the difference in the return to declining search frictions between two types depends on the difference in the extent of cross-sectional dispersion in their respective markets.

Our main empirical finding is to document a positive relationship between cross-sectional wage dispersion and growth in the data. We first show that cross-sectional wage dispersion is systematically lower in routine than nonroutine occupations. We then show that wage growth is systematically lower in occupations with less wage dispersion (routine) than in occupations with more wage dispersion (nonroutine). These findings are consistent with the hypothesis that declining search frictions have led to higher growth for nonroutine than for routine workers. In routine occupations, jobs are almost by definition all similar and, hence, a routine worker would have a similar productivity on different jobs (he is a jack of all trades). In nonroutine occupations, jobs are almost by definition differentiated and, hence, a nonroutine worker would be more productive in some jobs than in others (he is a master of one trade). Therefore, the return to declining search frictions would be lower for routine than nonroutine workers. A rudimentary calibration suggests that declining search frictions may account for 30 percent of the difference in growth between routine and nonroutine workers.

I. Environment

The economy is populated by a measure μ_i of workers of type $i = 1, 2, \ldots, N$, with $\mu_i \in [0,1]$ and $\sum_{i=1}^{N} \mu_i = 1$. Different types of workers belong to different labor markets. Within their own labor market, workers of type i are distributed uniformly along a circle with unit circumference. A worker's type captures characteristics (e.g., educational attainment, preferences over broad lines of work, etc…) which determine the market where the worker sells his labor. A worker's location along the circle captures the worker's skills. A worker maximizes the present value of income discounted at the rate $r > 0$. A worker's income is some endogenous wage $w_{i,t}$ when he is employed, and some exogenous income $b_{i,t}$ when he is unemployed.

The economy is also populated by a positive measure of firms. A firm maximizes the present value of profits discounted at the rate r. A firm hires workers by opening, maintaining, and filling vacancies. A firm opens a vacancy in labor market i and locates it randomly along the circle of unit circumference. The firm pays a flow cost $k_{i,t} > 0$ to maintain the vacancy. The location of the vacancy along the circle captures the skills required by the firm. Once the vacancy is filled, the firm and the hired worker produce an income flow of $y_{i,t} z_i(\epsilon)$, where $y_{i,t}$ is a component of productivity that is common to all matches between a firm and a worker of type i, and $z_i(\epsilon)$ is a component of productivity that is specific to a particular match between a firm and a worker of type i. Specifically, $z_i(\epsilon)$ is a strictly positive and decreasing function of the distance $\epsilon \in [0, 1/2]$ between the location of the worker and the location of the firm's job along the circle. We refer to ϵ as the firm-worker skill distance.

Unemployed workers and vacancies need to search the labor market in order to find each other.[2] We assume that the search activity of the two sides of the market generates a flow $A_t M(u_{i,t}, v_{i,t})$ of one-to-one contacts between unemployed workers and vacancies, $u_{i,t}$ and $v_{i,t}$ being the measures of jobless workers and vacant jobs in labor market i, $M(\cdot,\cdot)$ a CRS function, and A_t the efficiency of search. A searching worker contacts a vacancy at rate $A_t p(\theta_{i,t})$, where $\theta_{i,t} = v_{i,t}/u_{i,t}$ is the tightness of labor market i, $p(\theta) \equiv M(1,\theta)$ with the properties $p'(\cdot) > 0, p''(\cdot) < 0, p(0) = 0$, and $p(\infty) = \infty$. A vacancy contacts an unemployed worker at rate $A_t q(\theta_{i,t})$, where $q(\theta) \equiv p(\theta)/\theta$ with the properties $q'(\cdot) < 0$, $q(0) = \infty$, and $q(\infty) = 0$.

When a firm and a worker come into contact in labor market i, they observe the distance ϵ between them. Given ϵ, they choose whether to form a match. If the firm and the worker match, they negotiate the provisions of an employment arrangement. They then begin producing an income of $y_{i,t} z_i(\epsilon)$ and continue doing so until their match is terminated. If the firm and the worker do not match, the worker stays in the pool of unemployment and the firm's job stays in the pool of vacancies.

The Nash bargaining solution determines the provisions of the employment arrangement. Namely, it maximizes the product of the worker's and the firm's gains

[2] The results extend to a version of the model where workers are allowed to search on the job.

from trade taken, respectively, to the power of γ and $1-\gamma$, with $\gamma \in (0,1)$. The worker's gains from trade are defined as the value of the match to the worker net of the value of unemployment. The firm's gains from trade are defined, analogously, as the value of the match to the firm net of the value of an unfilled vacancy. The employment agreement specifies a path for the wage and, explicitly or implicitly, a time when the match ought to be terminated. We assume that employment agreements are sufficiently flexible to ensure that the termination time maximizes the sum of the value of the match to the firm and the worker (the joint value of the match). Since the wage transfers income from the firm to the worker at the rate of 1-to-1, the Nash bargaining solution is such that the worker and the firm get, respectively, a portion γ and $1-\gamma$ of the surplus, which is defined as the difference between the joint value of the match and the sum of the outside options.

The economic environment is nonstationary, as both the production and the search technologies improve over time. Specifically, the common component of productivity $y_{i,t}$ grows at the exogenous rate $g_{y,i} \geq 0$, and the efficiency of the search process A_t grows at the exogenous rate $g_A > 0$. Additionally, the unemployment income $b_{i,t}$ and the vacancy cost $k_{i,t}$ grow, respectively, at the rates $g_{b,i}$ and $g_{k,i}$.

The model is a version of MM with N types of workers, who do not interact either in the labor or in the product market.[3] Different types of workers differ along two crucial dimensions. First, different types use different production technologies, each associated with its own level of productivity $y_{i,t}$ and its own growth rate $g_{y,i}$.[4] For example, one type of worker may use a production technology that grows at a high rate. Another type of worker may use a production technology that grows at a low rate. Second, different types face a different relationship $z_i(\epsilon)$ between the component of productivity that is specific to a particular match and the distance ϵ between the worker's individual skill and the job's skill requirements. For example, one type of worker may face a function $z_i(\epsilon)$ with a very low elasticity of z with respect to ϵ. This type of worker is a "jack of all trades," as he is essentially equally productive doing any kind of job (within his labor market). Another type of worker may face a function $z_i(\epsilon)$ with a very high elasticity of z with respect to ϵ. This type of worker is a "master of one trade," as he is much more productive doing a job that requires skills similar to those that he possesses, rather than a job that requires skills different from the ones he has.

II. BGP

We focus on a BGP, an initial condition of the economy together with an equilibrium path along which—for each type of worker—the unemployment, vacancy, EU, and UE rates are time invariant. In a model with only one type of worker, the focus on a BGP is natural given the lack of any clear secular trend in the aggregate unemployment, vacancy, EU, and UE rates during the last 100 years of US history

[3] The results of the paper can be extended to an environment in which the output of different types is combined into a final good through a CES aggregator.

[4] We refer to $y_{i,t}$ as productivity and $g_{y,i}$ as the growth rate of the production technology. More generally though, $y_{i,t}$ represents income and depends on both technology (which affects the quantity of output) and preferences (which affect the price of the output). Similarly, $g_{y,i}$ is the growth rate of income, which depends on technological progress and preferences.

(see MM). In a model with multiple types of workers, the focus on a BGP seems like a natural starting point, as the stationarity of unemployment, vacancy, and transition rates for each type of worker guarantees the stationarity of these objects at the aggregate level.[5]

Let $V_{i,t}(z)$ denote the joint value of an employment relationship of quality z between a firm and a worker of type i. Let $U_{i,t}$ denote the value of unemployment for a worker of type i. Let $S_{i,t}(z) \equiv V_{i,t}(z) - U_{i,t}$ denote the surplus of a match of quality z between a firm and a worker of type i. Further, let $\theta_{i,t}$ denote the tightness of labor market i, $u_{i,t}$ the fraction of unemployed workers of type i, and $G_{i,t}(z)$ the c.d.f. of match qualities of employed workers of type i. The initial condition of the economy is the distribution of workers at date $t = 0$, that is, $\{u_{i,0}, G_{i,0}\}$. A path for $V_{i,t}$, $U_{i,t}$, $\theta_{i,t}$, $u_{i,t}$, and $G_{i,t}$ is an equilibrium if decisions are optimal, markets clear, and the dynamics of $u_{i,t}$ and $G_{i,t}$ are those implied by the agents' decisions and the initial condition. A BGP is an initial condition together with an equilibrium such that the unemployment, vacancy, and transition rates are time invariant for all types of workers, while the distribution $G_{i,t}$ grows at some constant endogenous rate $g_{z,i}$—in the sense that every quantile of the distribution grows at the rate $g_{z,i}$.

The joint value $V_{i,t}(z)$ of a z-quality match is

$$V_{i,t}(z) = \max_{T \geq 0} \int_t^{t+T} e^{-r(\tau - t)} y_{i,\tau} z \, d\tau + e^{-rT} U_{i,t+T}. \tag{1}$$

For any calendar time $\tau \in [t, t+T]$, the match between the firm and the worker generates income $y_{i,\tau} z$. At calendar time $t + T$, the match is terminated, at which point the worker's continuation value is $U_{i,t+T}$ and the firm's continuation value is the value of an unfilled vacancy, which equals zero because firms are free to open and close vacancies at will.

The optimal termination date T satisfies

$$z y_{i,t+T} + \mathring{U}_{i,t+T} \leq r U_{i,t+T}, \quad \text{and} \quad T \geq 0, \tag{2}$$

where one of the two weak inequalities must hold as an equality.[6] On the left of (2) is the gain from postponing the termination of the match by an instant, which is equal to the income of the match plus the time change in the value of unemployment. On the right of (2) is the cost of postponing the termination by an instant, which is equal to the annuitized value of unemployment. Overall, condition (2) states that either $T = 0$ and the gain from postponing the termination of the match by an instant is smaller than the cost, or $T > 0$ and the gain and cost are equal.

We define the reservation quality $R_{i,t}$ as

$$y_{i,t} R_{i,t} = r U_{i,t} - \mathring{U}_{i,t}. \tag{3}$$

It follows from (2) that an ongoing firm-worker match survives if and only if $z \geq R_{i,t}$. Moreover, a new firm-worker match is consummated if and only if $z \geq R_{i,t}$. That

[5] Moreover, in online Appendix C, we document that the low-frequency dynamics of unemployment, EU, and UE rates at the three-digit occupation level are similar to those of aggregate unemployment, EU, and UE rates.

[6] Condition (2) is necessary and sufficient (see MM).

is, $R_{i,t}$ is the minimal match-specific component of productivity such that an ongoing match survives and a new match is consummated. From (1) and (2), it also follows that $S_{i,t} = V_{it}(z) - U_{i,t}$ is strictly positive if and only if $z > R_{i,t}$.

The value $U_{i,t}$ of unemployment for a worker of type i is

$$rU_{i,t} = b_{i,t} + A_{i,t} p(\theta_i) \gamma \int_0^{z_i^{-1}(R_{i,t})} S_{i,t}(z_i(\epsilon)) 2\, d\epsilon + \mathring{U}_{i,t}. \tag{4}$$

The left-hand side of (4) is the annuitized value of unemployment. The right-hand side is the sum of three terms. The first is the worker's unemployment income. The second is the worker's option value of searching, which equals the rate at which the worker finds a vacancy multiplied by a portion γ of the expected surplus. The last term is the derivative of the value of unemployment with respect to calendar time t. Note that the formula for the expected surplus makes use of the fact that the skill distance ϵ is uniformly distributed over $[0, 1/2]$ and only matches with $\epsilon \leq z^{-1}(R_{i,t})$ are consummated.

The tightness of labor marker i satisfies

$$k_{i,t} = A_{i,t} q(\theta_i)(1-\gamma) \int_0^{z_i^{-1}(R_{i,t})} S_{i,t}(z_i(\epsilon)) 2\, d\epsilon. \tag{5}$$

The left-hand side of (5) is the cost to the firm from maintaining a vacancy in labor market i. The right-hand side is the gain, which equals the rate at which the vacancy finds a worker multiplied by a portion $1-\gamma$ of the expected surplus. Condition (5) states that θ_i equates the cost and gain of maintaining an additional vacancy in labor market i.

By definition of a BGP, the unemployment rate, the market tightness, and the EU and UE rates must be time-invariant for each worker type. These requirements are fulfilled if and only if

$$G'_{i,t}(R_{i,t}) \mathring{R}_{i,t} = h^i_{eu}, \tag{6}$$

$$A_t p(\theta_{i,t}) 2 z_i^{-1}(R_{i,t}) = h^i_{ue}, \tag{7}$$

$$u_i h^i_{ue} = (\mu_i - u_i) h^i_{eu}. \tag{8}$$

Condition (6) requires the EU rate to be time invariant. The EU rate at calendar time t is the density of $G_{i,t}$ at $R_{i,t}$ times the derivative of $R_{i,t}$ with respect to t. Condition (7) requires the UE rate to be time-invariant. The UE rate at calendar time t is the rate at which jobless workers find firms multiplied by the probability that their distance ϵ falls below the threshold $z^{-1}(R_{i,t})$. Condition (8) guarantees that the unemployment rate is time invariant by requiring that the flows of workers in and out of unemployment are equal. The condition for the time invariance of the market tightness is implicit in (5), which requires the market tightness to be constant.

By definition of a BGP, the distribution $G_{i,t}(z)$ must grow at some endogenous, constant rate $g_{z,i}$. That is, $z_{i,t}(x) = z_{i,0}\exp(g_{z,i}t)$ for all $x \in [0,1]$ and all $t \geq 0$, where $z_{i,t}(x)$ is the xth quantile of $G_{i,t}$. The condition is satisfied if and only if

$$(9) \qquad (\mu_i - u_{i,t})\, G'_{i,t}(z_{i,t}(x))\, z_{i,t}(x)\, g_{z,i} + u_i A_t p(\theta_i) 2\left[z^{-1}(R_{i,t}) - z^{-1}(z_{i,t}(x))\right]$$
$$= (\mu_i - u_{i,t})\, G'_{i,t}(R_{i,t})\, \mathring{R}_{i,t}.$$

The left-hand side of (9) is the worker flow into matches with a quality that is below an xth quantile that grows at the rate $g_{z,i}$. The other side of (9) is the worker flow out of matches with a quality that is below an xth quantile that grows at the rate $g_{z,i}$. Condition (9) thus ensures the time invariance of the fraction of workers who are in matches of a quality that is below an xth quantile that grows at the rate $g_{z,i}$.

III. Unequal Growth

For each worker type, the definition of a BGP is the same as in MM. Therefore, the conditions under which a BGP exists are exactly the same as in MM (Theorems 1 and 2) applied to each worker type. The conditions are of "if and only if" nature and are reported below, together with the characterization of the properties of a BGP.

THEOREM 1 (Existence and Properties of a BGP):

(*i*) *A BGP exists if and only if for each worker type $i = 1, 2, \ldots, N$ (a) the function $z_i(\epsilon)$ has the isoelastic form $\underline{z}_i \epsilon^{-1/\alpha_i}$ for some scale parameter $\underline{z}_i > 0$ and elasticity parameter $\alpha_i > 1$, (b) the growth rate $g_{b,i}$ of the unemployment income and the growth rate $g_{k,i}$ of the vacancy cost are given by $g_{y,i} + g_A/\alpha_i$, and (c) the discount rate r is greater than $g_{y,i} + g_A/\alpha_i$.*

(*ii*) *If the conditions above are satisfied, the BGP is unique and, for each worker type $i = 1, 2, \ldots, N$, such that (a) unemployment, vacancy, EU, and UE rates are time invariant; (b) $G_{i,t}$ is a Pareto distribution with coefficient α_i truncated at $R_{i,t}$, and grows at the constant rate $g_{z,i} = g_A/\alpha_i$; (c) average labor productivity grows at the rate $g_{y,i} + g_A/\alpha_i$.*

The first part of Theorem 1 lists "if and only if" conditions under which a BGP exists. The intuition is simple. On the one hand, improvements in the search technology increase the speed at which the unemployed find vacant jobs. On the other hand, they increase $R_{i,t}$, thus reducing the probability that a meeting is consummated. The two effects offset each other, and thus the UE rate is time invariant if and only if $z_i(\epsilon)$ has some constant elasticity $1/\alpha_i$ and the growth rates $g_{b,i}$ and $g_{k,i}$ equal $g_{y,i} + g_A/\alpha_i$. The EU rate is time invariant because the reservation quality $R_{i,t}$ grows at a constant rate and the distribution $G_{i,t}$ is Pareto. The unemployment rate is time invariant because so are the UE and EU rates.

The conditions on $g_{b,i}$ and $g_{k,i}$ may appear to be knife-edge, but they are not (see MM, appendix B). In fact, $g_{y,i} + g_A/\alpha_i$ is the growth rate of the average productivity of workers of type i. Thus, the cost of a vacancy grows endogenously at rate $g_{y,i} + g_A/\alpha_i$ as long as the input in maintaining a vacancy in labor market i is some constant amount of labor supplied by workers of type i. For the same reason, the unemployment income grows endogenously at rate $g_{y,i} + g_A/\alpha_i$ as long as $b_{i,t}$ is some constant fraction of the average output of workers of type i.

The second part of Theorem 1 lists the properties of a BGP. First, in a BGP, unemployment, vacancy, and transition rates are time invariant for each type of worker and, hence, are time invariant in the aggregate. Second, the distribution $G_{i,t}$ for employed workers of type i is Pareto with coefficient α_i truncated at the reservation quality $R_{i,t}$ and grows at rate $g_{z,i} = g_A/\alpha_i$. To understand this second property, note that when a firm and a worker of type i meet, the distance ϵ is a random draw from a uniform distribution with support $[0, 1/2]$. The match between the firm and the worker is consummated if and only if $\epsilon \leq z^{-1}(R_{i,t})$, and if so, the quality of the match is given by $z_i(\epsilon) = \underline{z}_i \epsilon^{-1/\alpha_i}$. Therefore, the distribution of quality between newly created matches is

$$\Pr(\tilde{z} \leq z) = \Pr\left(\epsilon \geq z^{-1}(z) | \epsilon \leq z^{-1}(R_{i,t})\right) \tag{10}$$

$$= 1 - \left(R_{i,t}/z\right)^{\alpha_i}.$$

The above expression shows that the quality distribution of new matches is a Pareto with coefficient α_i truncated at $R_{i,t}$. Since previously created matches are also distributed as a Pareto with coefficient α_i and survive if and only if their quality exceeds $R_{i,t}$, the overall distribution $G_{i,t}$ is given by (10) as well. The distribution $G_{i,t}$ grows at the rate $g_{z,i} = g_A/\alpha_i$, as this is the growth rate of $R_{i,t}$.

The second part of Theorem 1 also states that average productivity of labor for workers of type i grows at rate $g_{y,i} + g_A/\alpha_i$. To see this, note that $G_{i,t}$ is given by (10), and hence,

$$\int y_{i,t} z \, dG_{i,t}(z) = \frac{\alpha_i}{\alpha_i - 1} y_{i,t} R_{i,t}. \tag{11}$$

The above expression grows at the rate $g_{y,i} + g_A/\alpha_i$, as $y_{i,t}$ grows at rate $g_{y,i}$ and $R_{i,t}$ grows at rate $g_{z,i} = g_A/\alpha_i$. The growth rate of average productivity for workers of type i has two distinct sources. The first source is $g_{y,i}$: progress in the production technology used by workers of type i. The second source is g_A/α_i: progress in the search technology, which has a return on the productivity of workers of type i equal to the elasticity $1/\alpha_i$ of $z_i(\epsilon)$ with respect to ϵ. That is, the return of declining search frictions on the productivity of workers of type i is equal to the elasticity of the match-specific component of productivity with respect to the distance between the skills of a worker of type i and the requirements of a firm's job.

The growth rate of productivity translates into wage growth. Following Pissarides (1985) and Mortensen and Pissarides (1994), we assume that a worker and a firm bargain over the wage at time intervals of length $dt \to 0$. Then, the wage for a worker of type i in a match of quality z is given by

$$w_{i,t}(z) = \gamma y_{i,t} z + (1-\gamma) y_{i,t} R_{i,t}. \tag{12}$$

From (10) and (12), it follows that the wage distribution $L_{i,t}(w)$ for workers of type i is

$$L_{i,t}(w) = 1 - \left(\frac{\gamma y_{i,t} R_{i,t}}{w - (1-\gamma) y_{i,t} R_{i,t}}\right)^{\alpha_i}, \tag{13}$$

and the average wage for workers of type i is

$$\int w\, dL_{it}(w) = \left[\gamma \frac{\alpha_i}{\alpha_i - 1} + (1-\gamma)\right] y_{i,t} R_{i,t}. \tag{14}$$

Clearly, the average wage for workers of type i grows at the rate $g_{y,i} + g_A/\alpha_i$, which is the same as the growth rate of average productivity for that type of worker.

We are now in the position to state the main theorem of the paper.

THEOREM 2 (Unequal Growth): *Consider two worker types i and j, $i \neq j$. Wage growth for workers of type i may be higher than for workers of type j because (i) technological progress in production is biased toward workers of type i: $g_{y,i} > g_{y,j}$ and (ii) technological progress in search has a higher rate of return for workers of type i: $1/\alpha_i > 1/\alpha_j$.*

Theorem 2 identifies two sources of unequal wage growth for different groups of workers. The first source of unequal growth between workers of types i and j is that progress in the production technology is biased in favor of workers of type i. This is the canonical explanation for the rise in the college premium (that is, progress in the production technology is biased in favor of college graduates) or for the decline in the wages of routine workers relative to nonroutine workers (that is, progress in automation erodes the surplus generated by routine workers).

The second source of unequal growth is novel. Specifically, the second source of unequal growth between workers of types i and j is that progress in the search technology has a higher rate of return for workers of type i. The logic is simple. Suppose that workers of type i are masters of one trade (or specialists), in the sense that their productivity is more elastic to the distance between their idiosyncratic skills and the skill requirements of their job. In contrast, workers of type j are jacks of all trades (or generalists), in the sense that their productivity is less elastic to the distance between their idiosyncratic skills and the skill requirements of their job. Declining search frictions allow both workers of type i and workers of type j to become more selective with respect to the jobs they accept. The increase in selectivity, however, is going to have a higher rate of return in terms of productivity—and, hence, wages—for workers of type i than for workers of type j.

As a matter of interpretation, it is useful to point out that the notion of specialists and generalists is relative to jobs. That is, workers of type i may be more specialized than workers of type j because they are more productive at jobs that suit them well and less productive at jobs that do not suit them well. Equivalently, workers of type i may be more specialized than workers of type j because the jobs that are available to them are more heterogeneous in terms of skill requirements. In either case, workers of type i end up facing more heterogeneity in the match-specific component of their productivity.

Declining search frictions cause unequal growth if and only if different types of workers have a different elasticity of productivity with respect to the firm-worker skill distance. In turn, differences in the elasticity of productivity with respect to the firm-worker skill distance manifest themselves as differences in wage dispersion. Hence, declining search frictions cause unequal growth if and only if there are differences in the extent of wage dispersion across different types of workers. Specifically, declining search frictions cause higher growth for the types of workers displaying more wage dispersion.

It is easy to formalize the above argument. From (13), it follows that the ratio between the x_1th and the x_0th quantiles of the wage distribution for workers of type i is given by

$$\frac{w_i(x_1)}{w_i(x_0)} = \frac{\gamma(1-x_1)^{-1/\alpha_i} + (1-\gamma)}{\gamma(1-x_0)^{-1/\alpha_i} + (1-\gamma)}. \tag{15}$$

Abstracting from differences in bargaining power, the wage quantile ratio in (15) is different for two types of workers i and j if and only if the elasticity of productivity with respect to the firm-worker skill distance is different. In turn, the growth rate induced by declining search friction is different if and only if the elasticity is different. Moreover, since (15) is strictly increasing in $1/\alpha$, the growth rate g_A/α induced by declining search frictions is higher for workers of type i than for workers of type j if and only if the wage quantile ratio (15) is higher for workers of type i.

We have thus established the following result.

THEOREM 3 (Wage Dispersion and Growth): *Consider two groups of workers i and j,* $i \neq j$, *with* $g_{y,i} = g_{y,j}$. *Wage growth for workers of type i is higher than for workers of type j if and only if, for any* x_0, x_1 *in* $[0,1]$ *with* $x_1 > x_0$, *the ratio between the* x_1*th and the* x_0*th quantile of the wage distribution is higher for workers of type i than for workers of type j.*

IV. Empirics

Theorem 3 implies that, other things equal, there should be a positive relationship between wage dispersion and wage growth across different groups of workers. In this section, we document the existence of this relationship in the data and argue that declining search frictions can account for about 30 percent of the wage growth differential between routine and nonroutine workers. Specifically, we document that wage dispersion is systematically higher for workers in less routine occupations. We then document that wage growth for workers in occupations with more wage dispersion (less routine ones) is systematically higher than wage growth for workers in occupations with less wage dispersion (more routine ones). Lastly, we use our model to argue that some (but not all) of the differential wage growth between workers in routine and nonroutine occupations can be accounted for by differences in the return to declining search frictions.

We use wage and salary income data from the decennial census (1980) and the American Community Survey (2015). We restrict attention to workers aged 25 to 55 who are neither in the military nor enrolled in school. We further restrict attention to workers who work at least 35 hours per week, at least 48 weeks per year, and earn an annual income of at least $5,000 (measured in 2015 dollars). We divide annual income by total number of hours to compute hourly wages.

We proxy a worker's type with his three-digit occupation[7] using the crosswalk developed by David Dorn to maintain a consistent definition of occupations over time.[8] We measure wage dispersion in an occupation as the ratio between the seventy-fifth and the twenty-fifth percentiles of the wage distribution in 1980.[9] We measure wage growth in an occupation as the ratio between the average wage in 2015 and the average wage in 1980. Since differences in wage dispersion across occupations may be due to differences in the composition of workers with different observable characteristics, in online Appendix D, we replace wages with residuals from a Mincer regression of log-wages on gender, race, education, industry, and a quadratic polynomial of age. The results are qualitatively similar using raw or residual wages.

We classify each occupation based on the degree of routineness of the tasks that it involves, following Autor and Dorn (2013). We then group occupations into 20 equally sized clusters with increasing degree of routineness, so as to reduce noise. We define wage dispersion of a particular cluster as the average of the seventy-fifth to twenty-fifth wage percentile ratios in 1980 across all occupations that belong to that cluster. We define wage growth of a particular cluster as the average 2015–1980 wage growth across all occupations that belong to that cluster.

Figure 1, panel A shows that clusters of occupations that are more routine have a systematically lower wage dispersion than clusters of occupations than are less routine. The coefficient on routineness in an OLS regression of wage dispersion is -0.014 with a standard error of 0.0034 (significant at 1 percent) and the R^2 is 48 percent. This finding has a natural interpretation in light of our theory. In a routine occupation, jobs are—almost by definition—similar to each other. Hence, the productivity and wage of a worker with a particular set of skills is bound to be similar across different jobs in that occupation. In a nonroutine occupation, jobs are—again almost by definition—differentiated, and hence the productivity and wage of a worker with a particular set of skills may be quite different across different jobs in the same occupation. For this reason, a routine worker is a jack of all trades, in the sense that his productivity will be similar across different jobs. A nonroutine worker is a master of one trade, in the sense that his productivity may be quite higher in some jobs than in others.

Figure 1, panel B shows that clusters of occupations with higher wage dispersion have a systematically higher growth than clusters of occupations with lower wage dispersion.[10] The finding is consistent with our theory. Occupations with lower

[7] The choice of occupation as a proxy for a worker's type is not without flaws. In particular, while the type is assumed to be a permanent trait of a worker in the model, workers do switch occupations in the data. The discrepancy between model and data is somewhat mitigated by the fact that we group occupations with a similar degree of routineness.

[8] See online Appendix B for more details on the definition of an occupation and the measure of routineness.

[9] In online Appendix F, we use the ninetieth to tenth percentile ratio. The results are similar.

[10] Online Appendix E reports the analog of Figure 1, panel B without grouping occupations. The resulting relationship between wage dispersion and growth is still positive but noisier.

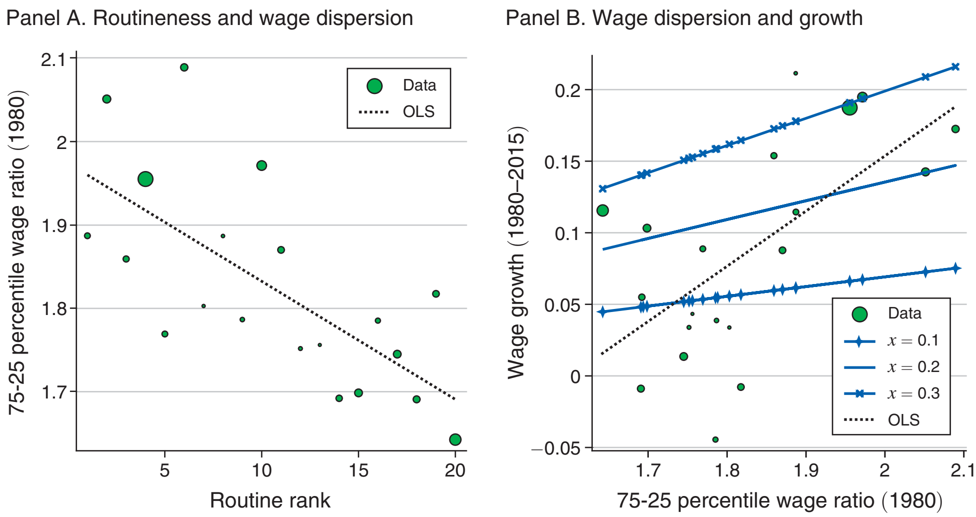

Figure 1. Wage Dispersion and Growth by Routineness

wage dispersion (which tend to be routine occupations) are such that the return to declining search frictions is lower, as workers in these occupations have similar productivity in all jobs. Occupations with higher wage dispersion (which tend to be nonroutine occupations) are such that the return to declining search frictions is higher, as the productivity of workers in these occupations varies more across jobs. The coefficient on wage dispersion in an OLS regression of wage growth is 0.384 with a standard error of 0.113 (significant at 1 percent), and the R^2 of the regression is 39 percent.

While the evidence presented in Figure 1, panel B is consistent with our theory, the slope of the OLS regression line may overstate the effect of wage dispersion on wage growth through differential returns to declining search frictions. In fact, it is quite possible that progress in the production technology is systematically higher in nonroutine occupations (which tend to have high wage dispersion) than in routine occupations (which tend to have low wage dispersion).

In order to tease out how much of the slope in the OLS regression line in Figure 1, panel B may be caused by differential returns to declining search frictions, we take a structural approach. First, using equation (15), we recover the rate of return to declining search frictions in the occupation cluster i, $1/\alpha_i$, from the observed wage dispersion in that cluster. Since some wage dispersion may be caused by heterogeneity in worker-specific productivity, we choose $1/\alpha_i$ to generate only some fraction x of wage dispersion in cluster i. Second, using the approach outlined in MM, we recover the rate g_A at which search frictions decline from the growth rate in the average number of applications received by a firm before filling its vacancy. We recover a g_A of about 2 percent per year between 1980 and 2010. Finally, by multiplying $1/\alpha_i$ and g_A, we recover the contribution of declining search frictions to the growth of wages in the occupation cluster i.

Figure 1, panel B illustrates our findings. The three lines represent the wage growth due to declining search frictions for $x = 0.1$ (bottom), 0.2 (middle), and 0.3

(top).[11] The three lines are drawn assuming that the worker's portion γ of the surplus is equal to 0.9.[12] For $x = 0.2$, the wage growth due to declining search frictions goes from about 9 percent for the most routine occupations (that is, those with the least wage dispersion) to about 15 percent for the least routine occupations (that is, those with the most wage dispersion). In the data, wage growth for the most routine occupations is about 3 percent and wage growth for the least routine occupations is about 20 percent. Thus, heterogeneous returns to declining search frictions account for about 6 out of 17 percentage points (about 30 percent) of the growth differential between routine and nonroutine occupations. The residual wage growth differential may be due to biased technological change, as suggested by Autor and Dorn (2013). If we assume that x is 0.1 (0.3) rather than 0.2, that is, we assume that the fraction of wage dispersion coming from differences in match-specific productivity is lower (higher), we find that declining search frictions generate lower (higher) growth for all occupations, and they account for about 3 (7) percentage points of the growth differential between routine and nonroutine occupations. In all cases, the message is similar: declining search frictions appear to have benefited nonroutine more than routine workers.

REFERENCES

Acemoglu, Daron, and Pascual Restrepo. 2018. "The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment." *American Economic Review* 108 (6): 1488–1542.

Albrecht, James W., and Bo Axell. 1984. "An Equilibrium Model of Search Unemployment." *Journal of Political Economy* 92 (5): 824–40.

Autor, David H., and Daivd Dorn. 2013. "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market." *American Economic Review* 103 (5): 1553–97.

Bagger, Jesper, François Fontaine, Fabien Postel-Vinay, and Jean-Marc Robin. 2014. "Tenure, Experience, Human Capital, and Wages: A Tractable Equilibrium Search Model of Wage Dynamics." *American Economic Review* 104 (6): 1551–96.

Bhuller, Manudeep, Andreas R. Kostøl, and Trond C. Vigtel. 2019. "How Broadband Internet Affects Labor Market Matching." http://conference.iza.org/conference_files/DATA_2019/Kostøl_a8219.pdf.

Bontemps, Christian, Jean-Marc Robin, and Gerard J. van den Berg. 2000. "Equilibrium Search with Continuous Productivity Dispersion: Theory and Nonparametric Estimation." *International Economic Review* 41 (2): 305–58.

Burdett, Kenneth, and Dale T. Mortensen. 1998. "Wage Differentials, Employer Size, and Unemployment." *International Economic Review* 39 (2): 257–73.

Card, David, and John E. DiNardo. 2002. "Skill-Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles." *Journal of Labor Economics* 20: 733–83.

Gregory, Victoria. 2019. "Firms as Learning Environments: Implications for Earnings Dynamics and Job Search." Unpublished.

Katz, Lawrence F., and Kevin M. Murphy. 1992. "Changes in Relative Wages, 1963–1987: Supply and Demand Factors." *Quarterly Journal of Economics* 107: 35–78.

Krusell, Per, Lee E. Ohanian, José-Víctor Ríos-Rull, and Giovanni L. Violante. 2000. "Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis." *Econometrica* 68 (5): 1029–53.

[11] For $x = 0.2$, the average α is 7. For $x = 0.1$, the average α is 12. For $x = 0.3$, the average α is 4.5. These are conservative values for the thickness $1/\alpha$ of the tail of the distribution of match-specific productivities. For instance, using data on the wage paths of individual workers and a model similar to ours, Martellini (2019) estimates α to be 3.6.

[12] For the sake of simplicity, our model abstracts from search on the job. In on-the-job search models, multiple firms compete for the services of a particular worker, and hence the share of the surplus accruing to workers ends up being quite large (see, e.g., Bagger et al. 2014; Menzio, Telyukova, and Visschers 2016; Gregory 2019 or Martellini 2019). We set γ equal to 0.9 to capture this important feature of on-the-job search models.

Lise, Jeremy, and Fabien Postel-Vinay. 2020. "Multidimensional Skills, Sorting, and Human Capital Accumulation." *American Economic Review* 110 (8): 2328–76.

Martellini, Paolo. 2019. "Local Labor Markets and Aggregate Productivity." Unpublished.

Martellini, Paolo, and Guido Menzio. 2020. "Declining Search Frictions, Unemployment and Growth." *Journal of Political Economy* 128 (12): 4387–4437.

Martellini, Paolo, and Guido Menzio. 2021. "Replication Data for: Jacks of All Trades and Masters of One: Declining Search Frictions and Unequal Growth." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. https://doi.org/10.3886/E129721V1.

Menzio, Guido, Irina A. Telyukova, and Ludo Visschers. 2016. "Directed Search over the Life Cycle." *Review of Economic Dynamics* 19: 38–62.

Mortensen, Dale T., and Christopher A. Pissarides. 1994. "Job Creation and Job Destruction in the Theory of Unemployment." *Review of Economic Studies* 61: 397–415.

Pissarides, C. A. 1984. "Efficient Job Rejection." *Economic Journal* 94: 97–108.

Pissarides, C. A. 1985. "Short-Run Equilibrium Dynamics of Unemployment, Vacancies, and Real Wages." *American Economic Review* 75 (4): 676–90.

Postel-Vinay, Fabien, and Jean-Marc Robin. 2002. "Equilibrium Wage Dispersion with Worker and Employer Heterogeneity." *Econometrica* 70 (6): 2295–2350.