Labor Market Returns to Personality: A Job Search Approach to Understanding Gender Gaps

Christopher Flinn, Petra E. Todd and Weilong Zhang *

Abstract

This paper investigates the effects of the Big Five personality traits on labor market outcomes and on gender disparities within a job search, matching and bargaining model with heterogeneous workers. In the model, parameters pertaining to productivity, job offer arrival rates, job dissolution rates and the division of surplus depend on worker education, cognitive skills, personality traits and other demographic characteristics. The model is estimated using a representative panel dataset, the German Socio-Economic Panel (GSOEP). Results show that both cognitive and noncognitive traits are important determinants of wage and employment outcomes. Higher levels of conscientiousness and emotional stability and lower levels of agreeableness increase hourly wages and promote greater job finding rates. A decomposition analysis shows that gender differences in two personality traits - agreeableness and emotional stability - primarily account for the gender wage gap and that their effect operates largely through a reduction in the bargaining power of women. Using results drawn from the clinical psychology literature, we find that a mental health intervention targeted at individuals with low levels of emotional stability could reduce the gender wage gap by from 2 to 5 percent and reduce wage inequality.

*Christopher Flinn is Professor of Economics at New York University and a Research Fellow at Collegio Carlo Alberto. Petra Todd is Edmund J. and Louise W. Kahn Term Professor of Economics at the University of Pennsylvania and a member of the U Penn Population Studies Group, NBER and IZA. Weilong Zhang is an assistant professor at University of Cambridge. An earlier version of this paper circulated under the title “Personality Traits, Job Search and the Gender Wage Gap.” We thank Magne Mogstad, Jose-Victor Rios-Rull, Holger Sieg, Richard Blundell, Carlos Carrillo Tudela, Kai Liu, Noriko Amano Patino, Tommas Trivieri, Tatiana Baron, and five anonymous referees for their very helpful comments. We also thank participants in conferences at Duke University and the University of Chicago and in seminars at the Chicago Federal Reserve Bank, Bocconi, University of Stockholm, Cambridge, Warwick, UVA, PSU, NYU, UPenn, U of Washington, and U of Miami for their comments.

1 Introduction

Despite substantial convergence in gender wage and employment differentials over the 1970s and 1980s, significant differences remain with women earning on average 25 percent less than men (Blau and Kahn (2006), Flabbi (2010b)). A large empirical literature uses data from the US and Europe to investigate the reasons for gender disparities. Individual attributes, such as years of education and work experience, account for a portion of gender wage and employment gaps, but a substantial unexplained portion remains. The early literature on gender wage gaps generally attributed residual gaps to unobserved productivity differences and/or labor market discrimination.

In recent decades, however, there is increasing recognition that non-cognitive skills, such as personality traits, are important determinants of worker productivity and may also contribute to gender disparities. The most commonly used noncognitive measurements are the so-called "Big Five" personality traits, which measure an individual's openness to experience, conscientiousness, extraversion, agreeableness and neuroticism (the opposite of emotional stability).[1] Figure 1 compares the distribution of the Big Five personality traits in our data for women and men. Women are more likely to score in the highest categories on openness to experience, conscientiousness, extraversion and agreeableness and in the lowest categories on emotional stability. Similar patterns have been documented for many countries and these trait differences have been shown to be significantly associated with gender wage gaps (e.g. Nyhus and Pons (2005), Heineck (2011), Mueller and Plug (2006), Braakmann (2009), Cattan (2013)). However, the mechanisms through which personality traits affect labor market outcomes have not been explored.

This paper examines the relationship between personality traits and labor market outcomes within a partial-equilibrium job search model. We develop and estimate a model in which personality traits potentially operate through multiple channels. In the model, workers who are heterogeneous in their characteristics stochastically receive employment opportunities from firms characterized in terms of idiosyncratic match productivity values. Firms and workers cooperatively divide the match surplus from the job using a predetermined worker's share parameter. We propose a new way of incorporating individual heterogeneity into the search framework by specifying job search parameters as index functions of a possibly high-dimensional set of attributes, including both cognitive and noncognitive individual characteristics. We use the model to explore how cognitive and noncognitive traits affect hourly wages, employment and labor market dynamics and to better understand the

[1]The measures aim to capture patterns of thoughts, feelings and behavior that correspond to individual differences in how people actually think, feel and act (Borghans et al. (2008), Almlund et al. (2011)).

Figure 1: The distribution of Big Five personality traits by gender

(a) Openness to experience

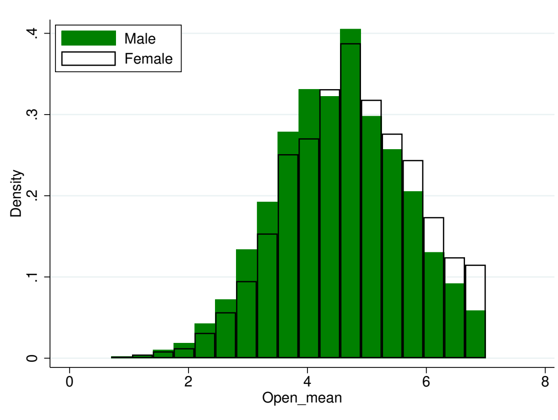

(b) Conscientiousness

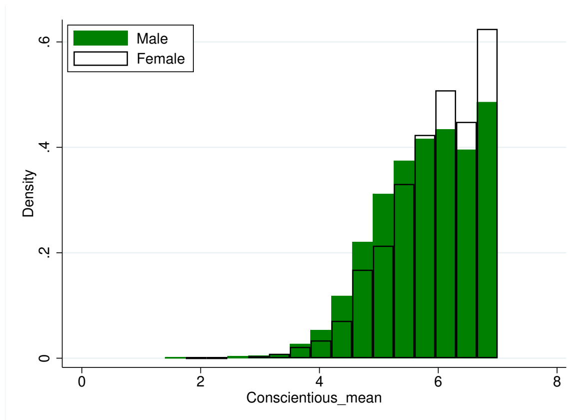

(c) Extraversion

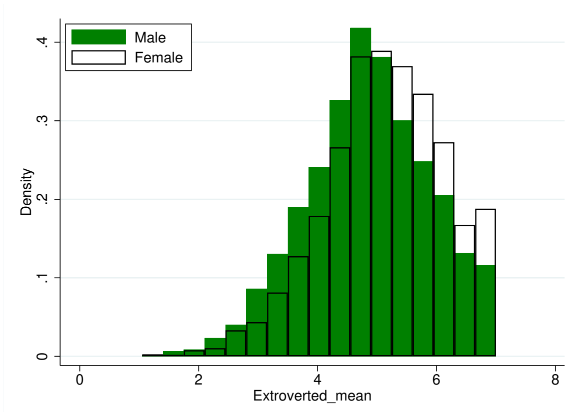

(d) Agreeableness

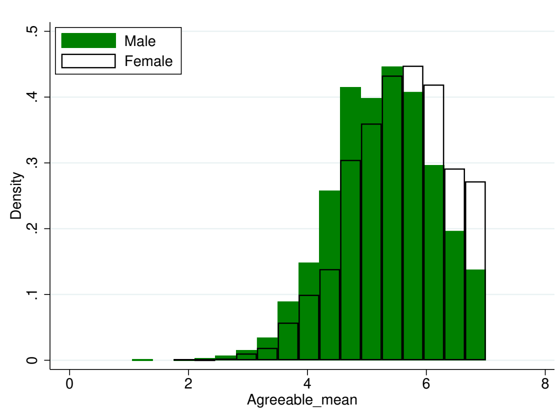

(e) Emotional Stability

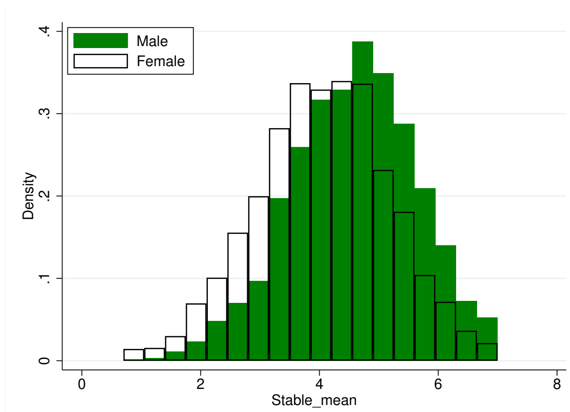

Notes: The measures are based on the scores of individuals aged 25 to 60 who report personality traits in the GSOEP. Each trait measure is defined on a scale of 1 to 7.

determinants of gender wage gaps. Understanding the mechanisms through which gender gaps arise is important for designing effective policies to address gender disparities in labor market outcomes.

Our model builds on traditional matching-bargaining models, such as Cahuc et al. (2006) and Dey and Flinn (2005). It also contributes to a smaller literature that uses job search models to analyze gender wage gaps (e.g. Bowlus and Grogan (2008), Flabbi (2010a), Liu (2016), Morchio and Moser (2020), Xiao (2020), Amano-Patino et al. (2020)). Our modeling approach differs from prior studies by allowing job search parameters to depend, in a flexible way, on a larger set of worker characteristics including personality traits. We quantify the importance of workers' characteristics operating through four distinct channels: worker productivity, job finding rates, job exit rates, and bargaining power.[2]

Model parameters are estimated by maximum likelihood using data from the German Socioeconomic Panel (GSOEP), a large, representative, longitudinal sample of German households. We focus on working age (age 25-60) individuals surveyed in 2013 and followed until 2019 (the most recent year of data available). We use information on their gender, age, education, cognitive skills, wages, job transitions, and on the Big Five personality trait measurements. We show that personality traits are significantly associated with hourly wages and unemployment/employment spell lengths.

We estimate two different, but nested, job search models that incorporate varying degrees of individual heterogeneity. In the most general specification, worker productivity, job arrival rates, job exit rates, and bargaining parameters all depend, through indices, on a comprehensive set of worker characteristics that include cognitive and non-cognitive skill measures. In the more restricted specification, we allow parameters to only vary by gender. Likelihood ratio tests overwhelmingly reject the restrictive specification in favor of the more flexible one, which also clearly provides a better visual fit to the data. As is to be expected, the specification that incorporates more observable dimensions of individual heterogeneity assigns a lesser role to unobserved model components (e.g. match quality, measurement error) in fitting the data.

Using our estimated heterogeneous job search model, we simulate steady state labor market outcomes for men and women. We analyze how each of the cognitive traits (education, cognitive skills) and each of the personality traits, ceteris paribus, influences labor market outcomes. We find that the effects of personality traits on men and women are qualitatively similar but quantitatively different. For example, a one standard deviation increase in con-

[2] In the estimation of structural search models, conditioning variables are often used to define labor markets, and then estimation proceeds as if these labor markets are isolated from one another. In our case, the labor market parameters are allowed to depend on a linear index of individual characteristics, which include personality measures and other individual characteristics.

scientiousness increases average wages by 6.8 percent for men and 5.3 percent for women. For both genders, conscientiousness and emotional stability increase hourly wages and shorten unemployment spells, whereas agreeableness leads to worse labor market outcomes.

In order to assess the relative importance of personality traits and other characteristics in explaining gender wage gaps, we perform a decomposition similar in spirit to a Oaxaca-Blinder decomposition but adapted to our nonlinear setting. We find that gender differences in education and cognitive skills, in terms of levels or associated returns, cannot account for observed wage gaps. Instead, differences in personality traits emerge as the primary sources of wage gaps. Detailed investigation shows that agreeableness and emotional stability contribute the most to accounting for gender differences in wages. In particular, women's higher average levels of agreeableness and lower levels of emotional stability relative to men substantially reduce their bargaining power, resulting in lower wages.

Lastly, we use the estimated model to explore the effects of potential interventions aimed at modifying some aspects of personality. Most personality psychologists believe personality traits are relatively stable during adulthood (e.g. Costa Jr and McCrae (1988); McCrae and Costa Jr (1994)) and are not that responsive to common life events or on-going life experience (e.g. Lüdtke et al. (2011); Cobb-Clark and Schurer (2013, 2012); Bleidorn et al. (2018)), However, recent research in clinical psychology shows that there are some mental health interventions that are effective in changing some aspects of personality. (e.g. Barlow et al. (2014); Bagby et al. (2008); Soskin et al. (2012)). Roberts et al. (2017) perform a meta-analysis of the results from 207 studies in the clinical psychology literature and conclude that personality traits are modifiable with short-term (6 to 8 weeks) therapeutic treatments, with changes in emotional stability being the primary trait affected by the treatment. We consider the effects of targeting such interventions to individuals (both males and females) with low measured levels of emotional stability and find that the gender wage gap reduces by 2 to 6 percent, depending on the fraction of the population targeted.

Our evidence contributes to the literature analyzing gender differences in job search behaviors and outcomes. Most prior studies estimate different search parameters by gender and education groups (e.g. Bowlus (1997), Bowlus and Grogan (2008), Flabbi (2010a), Liu (2016), Morchio and Moser (2020), Amano-Patino et al. (2020)). In comparison, we allow job search model parameters to depend on a larger set of worker characteristics to account for both cognitive and non-cognitive dimensions of worker heterogeneity. There are two studies that empirically investigate the association between non-cognitive traits and job search, Caliendo et al. (2015) and McGee (2015). The non-cognitive measure used in both papers is "locus of control"(LOC), which is a measure of how much individuals think success depends on "internal factors"(i.e. their own actions) verses "external factors."[3] To the best of our knowledge, this is the first study to incorporate the Big Five personality traits into a job search, matching, and bargaining framework.

This paper also builds on a literature examining how personality traits relate to wages and employment. Many studies demonstrate that gender differences in personality traits are significantly associated with differences in wages and employment (e.g. Nyhus and Pons (2005), Heineck (2011), Mueller and Plug (2006), Braakmann (2009), Cattan (2013)). A few studies investigate the relationship between personality traits and gender wage gaps by adopting a standard Oaxaca-Blinder decomposition framework (e.g. Mueller and Plug (2006); Braakmann (2009); Nyhus and Pons (2012); Risse et al. (2018); Collischon (2021)). Some of these studies find that agreeableness and emotional stability importantly contribute to gender wage gaps. By incorporating personality traits into a canonical job search and bargaining model, our results provide further evidence for which cognitive and noncogitive traits matter the most for gender disparities and also quantify the main mechanisms behind them. For example, we find that the most important channel through which personality traits affect gender gaps is wage bargaining, rather than productivity or job search behavior. Our paper also contributes to a small literature incorporating personality traits into specific behavioral models. (Todd and Zhang (2020); Heckman and Raut (2016); Flinn et al. (2018)).

There are several studies in the workplace bargaining literature showing that women are less likely to ask for fair wages, both from lab experiments (e.g. Stuhlmacher and Walters (1999); Dittrich et al. (2014)) and survey data (e.g. Säve-Söderbergh (2007); Card et al. (2015)). However, there is no consensus on the reason for this phenomenon. Possible explanations offered include gender differences in risk preferences (e.g. Croson and Gneezy (2009)), attitudes towards competition (e.g. Lavy (2013); Manning and Saidi (2010)) and negotiation skills (e.g. Babcock et al. (2003)). Our results suggest that gender differences in personality traits are also a key factor underlying differences in bargaining outcomes. Specifically, we find that women's higher levels of agreeableness and lower levels of emotional stability relative to men reduce women's relative bargaining power. This finding is consistent with Evdokimov and Rahman (2014), who design a bargaining experiment and show that an increase in a worker's agreeableness level leads a manager to allocate less money to the worker.

This paper proceeds as follows. The next section presents our baseline job search model. In Section 3 we describe the data used in the estimation of the model. Section 4 contains a

[3] Previous studies generally indicate that higher internal LOC is positively correlated with earnings. However, LOC is not that relevant for gender wage gaps either in terms of differential endowments or returns. (see e.g. Semykina and Linz (2007); Heineck and Anger (2010); Nyhus and Pons (2012))

discussion of the econometric implementation of the model. Section 5 presents the parameter estimates of the model. In Section 6 we interpret the estimates and present wage decomposition results. In this section we also consider the impact of a hypothetical experiment in which individuals with low levels of emotional stability are given pyschotherapy on wage inequality and the gender wage gap. Section 7 concludes.

2 Model

We now introduce the canonical job search, matching, and bargaining model and then discuss how we extend the model to incorporate heterogeneous primitive parameters that are functions of observable individual attributes.

2.1 Setup and environment

The model is set in continuous time, with a continuum of risk-neutral and infinitely lived agents: firms and workers. Workers are distinguished by different observable "types," denoted by the vector z. An unemployed worker meets firms at the rate $\lambda_U(z)$, and an employed worker meets new potential employers at the rate $\lambda_E(z)$, where both of these rates are assumed to be exogenously determined.[4] An individual's general productive ability, which is constant over their labor market career, is denoted $a(z)$. When a worker matches with a firm, the match quality value θ is determined by a draw from the cumulative distribution $G_z(\theta)$, which has a corresponding probability density function $g_z(\theta)$, both of which are defined on R_+. The productivity (total surplus) when an individual z matches with a job with match quality value θ is given by [5]

$$y(z) = a(z) \times \theta$$

The flow value of unemployment to the individual of type z is assumed to be $a(z) \times b$, where b can differ by gender.[6] Employment matches are dissolved at the exogenous rate

[4]The exogeneity assumption regarding worker-firm contact rates is what makes our analysis "partial equilibrium." A general equilibrium version of the model would endogenize these rates.

[5]Postel-Vinay and Robin (2002) and Cahuc et al. (2006) use a similar functional form for the flow productivity $y = a\theta$, where a and θ denotes the worker's and firm's productivity type, receptively. Although the specification looks similar, the interpretation of θ is different, which is mainly driven by the nature of the data. In the case of Postel-Vinay and Robin (2002) and Cahuc et al. (2006), matched worker-firm information is available, enabling the authors to identify distributions of worker and firm types nonparametrically. To the best of our knowledge, there are no such data sets that report worker's personality traits. Therefore, our model must rely only on supply side data, but we do allow workers with different values of z (e.g. men and women) to have different match quality distributions.

[6]The assumption that the flow value of being unemployed is proportional to ability a is common in the literature (e.g. Postel-Vinay and Robin (2002); Bartolucci (2013); Flinn and Mullins (2015)). This assumption is critical for separately identifying the ability a distribution from distribution of match-specific

$\eta(z)$. The common discount rate of all agents in the model, firms and workers, is ρ, which is independent of z.[7] The worker and the firm bargain over the wage w using a protocol resembling Nash bargaining, with the outside option of the individual dependent upon the particular bargaining protocol assumed. The worker's flow payoff from this match is w and the firm's flow revenue is $y(z) - w$. The "bargaining power" of the individual is denoted $\alpha(z)$.

In our application, the type z of an individual corresponds to a vector of observable individual characteristics that includes gender, education level, cognitive skills, birth cohort, and the Big Five personality trait assessments.[8] Because the model is stationary, we assume that these characteristics are time-invariant.

2.2 Job search and wage determination

In order to reduce notational clutter, for now we suppress the notation that conditions each of the model parameters on z. We will reintroduce z below when we discuss the model's econometric implementation.

2.2.1 On-the-job search

Employed workers receive job offers from other employers at the rate λ_E. When an employee meets a new employer, the potential match quality value θ' is immediately revealed to the worker and the firm. The potential employer's initial wage offer and whether the employee leaves for the new job depends critically on whether the current employer is allowed to renegotiate the wage offer. There are two alternative assumptions that have been used in the literature. In Dey and Flinn (2005) and Cahuc et al. (2006), firms are able to observe the worker's productivity at the competing firm and the firms behave as Bertrand competitors, with the culmination of the bidding process resulting in the worker going to the firm where their match quality value is largest.[9] Bertrand competition is possible only if the incumbent firms can verify a worker's outside offers and commit to any counteroffers that they make.[10]

productivity, as is discussed in Section 4.2 below.

[7]There is some evidence that workers with different cognitive and non-cognitive ability tend to have different discount rates (Dohmen et al. (2011)). However, we do not allow for such dependence, because the (ρ, b) are not individually identified in the canonical search framework. See the discussion in Section 4.

[8]We incorporate gender as a state variable in a more flexible way than other observed characteristics. We will discuss the econometric specification in Section 2.3.

[9]It is important to note that this decision does not depend on the worker's time-invariant general ability a, since a is the same at all potential employers.

[10]The firm has an incentive to renege on its offered wage once the potential competitor's offer has been withdrawn. This is the case if wage offers at alternative employers are withdrawn as soon as they have been rejected by the searcher.

If these assumptions do not hold and current employers do not renegotiate wage offers with their employees, then the worker's outside option is the value of being unemployed. This "no renegotiation" case is considered in Flinn et al. (2017), for example. It is worth noting that both specifications generate efficient mobility, as workers only leave a current employer if the match productivity at the new employer is at least as great as current match quality value ($\theta' \geq \theta$). However, the models imply different wage processes over the labor market career.

Although assuming Bertrand competition between firms is theoretically appealing, it is not clear how realistic this assumption is. It can be shown that introducing a positive cost of negotiation discourages Bertrand competition and can make it unprofitable for firms to poach workers from other firms with better match values. Mortensen (2005) argues that counteroffers are empirically uncommon.[11] Moscarini (2008) claims moral hazard concerns largely explain why firms do not match outside offers. An implication of firms renegotiating wages to retain workers is that the wage at an employer can only increase over time. In contrast, no renegotiation implies that the wage is constant over the course of a job spell. Given some of the theoretical objections to the assumptions underpinning Bertrand competition between firms and considering the fact that the average real wage growth observed in our data at a continuous job spell is close to zero, we implement the model without renegotiation. For this reason, the outside option for wage determination at any potential job is the value of unemployed search.[12]

2.2.2 Worker and firm value functions

We now describe the value function $V_E(\theta; w, a)$ for an employed worker with individual productive ability a working at a firm where their match quality value is θ and their wage is w:

$$\rho V_E(\theta; w, a) = w + \underbrace{\eta \left(V_U(a) - V_E(\theta; w, a)\right)}_{(1)} + \lambda_E \underbrace{\int_\theta \left(V_E(x; a) - V_E(\theta; w, a)\right) dG(x)}_{(2)}, \tag{1}$$

[11]Mortensen (2003, p. 99) writes: "Unlike in the market for academic economists in the United States, making counteroffers is not the norm in many labor markets. More typically, a worker who informs his employer of a more lucrative outside option is first congratulated and then asked to clear out immediately."

[12]In a previous version of this paper in which we used a different German data set, we estimated both the models with and without renegotiation and found that the model without renegotation provided a much better fit to the data.

where $V_U(a)$ is the value of unemployment to an individual of ability a and ρ is the discount rate.[13] Term (1) corresponds to the case in which the current job with wage w is dissolved due to an exogenous shock, which occurs at rate η. Term (2) corresponds to the case in which the individual moves to a new firm with a better match value, $x > \theta$. $V_E(x; a)$ is the value of being employed at an alternative firm with a match quality value x for an individual of ability a at a wage that solves the bargaining problem specified below (equation 7).

The value of this match to the firm (given the same θ, a and w) is given by

$$\rho V_F(\theta; w, a) = a\theta - w + \underbrace{\eta\left(0 - V_F(\theta; w, a)\right)}_{(1)} + \underbrace{\lambda_E \int_\theta \left(0 - V_F(\theta; w, a)\right) dG(x)}_{(2)}$$

$$(2) \qquad \Rightarrow V_F(\theta; w, a) = \frac{a\theta - w}{(\rho + \eta + \lambda_E(1 - G(\theta)))}.$$

The term $a\theta$ is the flow revenue to the firm and $a\theta - w$ is the firm's flow profit. The current match can be dissolved either due to an exogenous shock (term (1)) or due to the worker moving to a different firm at which their match-specific productivity is higher (term (2)). When the worker leaves, the vacancy has a value equal to 0 (a normalization).

We assume that every job searcher with a general productive ability a has a flow utility in the unemployment state equal to ab, where b is a constant. The value of being unemployed for an individual with ability a is given by $V_U(a)$, where

$$(3) \qquad \rho V_U(a) = ab + \lambda_U \int_{\theta^*(a)} \left(V_E(x; a) - V_U(a)\right) dG(x)$$

where $\theta^*(a)$ is the reservation match value at which an individual with ability a is indifferent between employment and continued search, defined by

$$V_U(a) = V_E(\theta^*; a)$$

Let $V_U \equiv V_U(1)$ and $V_E(\theta) \equiv V_E(\theta; 1)$ denote the value functions for individuals with "standard" ability $a = 1$. It is straightforward to show that, for any $a > 0$, $V_U(a) = aV_U$ and $V_E(\theta; a) = aV_E(\theta)$.[14] Using these equalities and substituting $aV_U(a) = aV_E(\theta^*)$ into equation

[13] It is worth noting that θ is the only time-varying state variable. The individual's general productive ability a is considered to be time-invariant, but we include it as an explicit argument for purposes of deriving an important factorization result that follows.

[14] To see why, note that ability, a, enters proportionally in the flow value for unemployment (ab) and into the flow productivity for employment ($a\theta$). There are a number of papers in the literature that assume that

1, we obtain the following reservation wage equation:

$$w^*(a) = a\theta^*(a) = a\rho V_U - \lambda_E \int_{\theta^*(a)} (aV_E(x) - aV_U)\, dG(x) \tag{4}$$

$$\Rightarrow \theta^*(a) = \rho V_U - \lambda_E \int_{\theta^*(a)} (V_E(x) - V_U)\, dG(x) \tag{5}$$

Equation 5 highlights the reservation match value $\theta^*(a)$ is a constant value across different ability a (as the solution to $\theta^*(a)$ in equation 5 is independent of a). That is, $\theta^*(a) = \theta^*$ for all a. Therefore, the reservation wage $w^*(a)$ for an individual of ability a is simply the "common" shared match reservation wage θ^* multiplied by a. This implies that whether a particular job is accepted is only a function of the match quality at that job and not a. We derive these decision rules in more detail and discuss the computation of the reservation match value θ^* in Appendix A.1.1.

2.2.3 Wage setting

The wage for an individual of productive ability a at a firm at which their match productivity value is θ is the solution to the following maximization problem:

$$w(\theta; a) = \arg\max_w (V_E(\theta; w, a) - V_U(a))^{\alpha} V_F(\theta; w, a)^{1-\alpha} \tag{6}$$

where $V_E(\theta; w, a)$ is the value of employment to the productive ability a worker (given in equation 1) and $V_F(\theta; w, a)$ is the value to the firm (given in equation 2). The parameter α is the surplus division parameter, capturing the worker's share of the total surplus. One caveat is that although the wage determination problem is similar to the Nash bargaining protocol, the axioms of Nash bargaining do not hold in this case.[15] Therefore, the surplus division parameter should not be interpreted as Nash bargaining power parameter, but instead simply as a surplus division parameter. By inserting the wage determined in (6) into the value function $V_E(\theta; w, a)$, we obtain the value function $V_E(\theta; a) \equiv V_E(\theta; w(\theta; a), a)$,

the flow utility is proportional in a, e.g. Postel-Vinay and Robin (2002), Cahuc et al. (2006) and Flinn and Mullins (2015).

[15] Shimer (2006) argues that in a simple search-matching model with on-the-job search, the standard axiomatic Nash bargaining solution is inapplicable, because the set of feasible payoffs is not convex. This non-convexity arises because an increase in the wage has a direct negative effect on the firm's rents but also an indirect positive effect raising the duration of the job. Gottfries (2018) extends Shimer's model by allowing the renegotiation to occur scholastically at a Poisson rate γ and shows the Nash bargaining solution is justified in such as a model. Our specification is nested as the limiting case in which wages are never renegotiated ($\gamma \to 0$).

defined as the value of employment when the wage is equal to the wage defined in (6).

Because $V_E(\theta; w, a) = aV_E(\theta; w)$, $V_U(a) = aV_U$, and $V_F(\theta; w, a) = aV_F(\theta; w)$, we can rewrite the surplus division problem as

$$w(\theta; a) = a \arg\max_{w} (V_E(\theta; w) - V_U)^{\alpha} V_F(\theta; w)^{1-\alpha} \tag{7}$$

Then we can think of the surplus division problem as solving for the wage for an individual of "standard" ability $a = 1$. The solution of the bargained wage is also proportional to a. For an individual with a productivity level $a \neq 1$, the wage is given by $w(\theta; a) = aw(\theta)$, where $w(\theta) \equiv w(\theta; 1)$ (the bargained wage for an individual of ability $a = 1$). Its solution, solving from equation 7, is given by the following expression:

$$\begin{aligned} w(\theta; a) &= a\left[\alpha\theta + (1-\alpha)\left(\rho V_U - \lambda_E \int_{\theta^*} \left(V_E(x) - V_U\right) dG(x)\right)\right] \\ &= a[\alpha\theta + (1-\alpha)\theta^*] \\ &= a[\theta^* + \alpha(\theta - \theta^*)] \end{aligned} \tag{8}$$

This equation shows why α represents the worker's share of the total surplus. For every additional unit of match quality value, the worker gets share α while the firm gets the share $1-\alpha$. As the worker's share approaches 0, we have $w = a\theta^*$ so that the workers are indifferent between employment and unemployment. If instead $\alpha = 1$, we have $w = a\theta$ and the firm has a flow profit of 0 in this case.

From equation 8, we obtain $\theta = \frac{\frac{w(\theta;a)}{a} - (1-\alpha)\theta^*}{\alpha}$. It is straightforward to see this is an affine mapping from the left-truncated match distribution $G(\theta|\theta \geq \theta^*)$ into the bargained wage distribution $F(w|w \geq w^*)$, where the lowest accepted wage is equal to the productivity associated with the lowest accepted match value, $w^* = a\theta^*$. Then the cumulative distribution function of wages for an individual with ability a is

$$F(w) = \frac{G\left(\alpha^{-1}\left(\frac{w}{a} - (1-\alpha)\theta^*\right)\right) - G(\theta^*)}{1 - G(\theta^*)}, w \geq a\theta^* \tag{9}$$

and the corresponding conditional wage density is given by

$$f(w) = \frac{1}{\alpha a} \frac{g(\alpha^{-1}\left(\frac{w}{a} - (1-\alpha)\theta^*\right))}{1 - G(\theta^*)}, w \geq a\theta^*$$

We will use this derived bargained wage distribution $F(w)$ when constructing the likelihood function in Section 4.3.

2.2.4 Household search

In Flinn et al. (2018), we develop and estimate a static model of household bargaining over time allocation decisions using Australian data.[16] Because men and women often inhabit households together, their labor supply decisions can be thought of as being jointly determined. Gender differences in wages may reflect patterns of assortative mating in the marriage market as well as the manner in which household decisions are made.

In this paper, the linear flow utility assumption allows us to reconcile our model with a household model.[17] Both men and women are assumed to have flow utility functions given by their respective wages w when employed and by the constants ba when they are unemployed. The linear utility assumption allows the the household's maximization problem to be decentralized as the sum of two individual maximization problems as noted in Dey and Flinn (2008). Under this admittedly strong, albeit common, assumption, we do not have to jointly model the labor market decisions of household members.[18]

2.3 Incorporating individual heterogeneity

So far, we described the search and bargaining model given a set of labor market parameters $\Omega = \{\lambda_U, \lambda_E, \eta, \alpha, a, b, \sigma_\theta\}$.[19] We now reintroduce individual types z and describe how we allow search parameters to depend on worker characteristics, that include education, cognitive skills, personality traits, birth cohort and gender. Because the linear index functions associated with each of the parameters potentially take values anywhere in $\mathbb{R}$, and many of the parameters are defined on subsets of $\mathbb{R}$, we choose standard functions to map the index function values into the appropriate parameter space. For all of the primitive parameters defined on $\mathbb{R}_+$ we use the $\exp(\cdot)$ function and for the one primitive parameter (α) which is defined on $[0, 1]$ we use the logit transformation. While somewhat arbitrary, these are

[16]The model is a static household decision-making model in which husbands and wives decide on time allocation (numbers of hours) to allocate to market work, home production, or leisure. Households can be cooperative or noncooperative and we estimate the fraction that is cooperative. Household bargaining weights and wage offer equations can depend on personality traits. In this paper, our focus is on workers bargaining with firms rather than husbands and wives bargaining with each other. It would be interesting, though, to develop and estimate a model that includes both within family and worker-firm bargaining.

[17]This assumption is made also because it obviates the need to include a specification of the capital markets within which individuals operate, since there is no borrowing or saving under the risk neutrality assumption.

[18]Under the alternative assumption of nonlinear utility, bargaining between spouses as well as with firms must be considered. Even if capital markets were not introduced into the model, this would considerably complicate the analysis.

[19]The parameter σ_θ denotes the standard deviation of distribution of $\log \theta$, which is assumed to be normal (so that θ follows a lognormal distribution). We further assume $\mu_\theta = -0.5\sigma_\theta^2$ so that $E(\theta) = exp(\mu_\theta + 0.5\sigma_\theta^2) = 1$.

standard transformations used in a variety of estimation contexts. Then we have:

$$\Omega_i \equiv \begin{bmatrix} \alpha(i) \\ \eta(i) \\ a(i) \\ \lambda_U(i) \\ \lambda_E(i) \\ b(i), \sigma_\theta(i) \end{bmatrix} = \begin{bmatrix} \frac{\exp(z_i' \gamma_\alpha^g)}{1+\exp(z_i' \gamma_\alpha^g)} \\ \exp(z_i' \gamma_\eta^g) \\ \exp(z_i' \gamma_a^g) \\ \exp(z_i' \gamma_U^g) \\ \tau^g \lambda_U(i) \\ \text{differ only by gender} \end{bmatrix}, \tag{10}$$

where z_i include all observed heterogeneity: education level, cognitive ability, personality traits, and birth cohorts, except for gender of the individual. The γ_j^g are gender-specific index coefficients, where $g \in \{M, F\}$ denotes the gender and j refers to a particular primitive parameter. The gender-specific coefficients γ_j^g allow for potential asymmetries in how traits of men and women are valued in the labor market.

As indicated above, we assume the parameters $\{\alpha(i), \eta(i), a(i), \lambda_U(i)\}$ are all functions of z_i. For the job arrival rate $\lambda_E(i)$, we assume it to be proportional to $\lambda_U(i)$, $\lambda_E(i) = \tau^g \lambda_U(i)$, where the proportionality factor τ^g may be gender-specific.[20] We also assume that $b(i)$ and $\sigma_\theta(i)$ differ only by gender. Both of these terms enter into our model as a product with ability $a(i)$ (ab and $a\theta$), and $a(i)$ is a function of z_i.

3 The German socio-economic panel (GSOEP)

Our study uses the German Socio-Economic Panel (GSOEP), which is a large-scale representative longitudinal household survey. Every year, there were nearly 11,000 households surveyed and more than 20,000 persons sampled from the German residential population. We focus on individuals surveyed in 2013 and followed until 2019 (the latest available year). We exclude individuals younger than 25 or older than 60, because we do not model schooling decisions or retirement. The GSOEP collects core labor market outcomes in all waves. It also collects an individual's self-reported personality traits and cognitive abilities in selected years. Personality traits are usually considered to be fairly stable after age 30 (McCrae et al. (2000)). Some studies find that personality traits change somewhat over the life course but observe that the rate of change is modest, allowing for meaningful comparisons across

[20]It is worth noting that τ^g could be generalized to be a function of all characteristics z_i and the coefficients would still be identified in theory, and in fact, we have successfully estimated this less restrictive model. However, we do not have strong reasons to believe that the personality traits will affect $\lambda_E(i)$ and $\lambda_U(i)$ in different ways. In order to economize on the number of parameters estimated, we have imposed this restriction.

individuals.[21]

Personality traits. The Big Five personality traits are measured using a 15-item self-assessment short version of the Big Five Inventory (see Appendix Table A2). Compared to the most widely used revised NEO Personality Inventory (NEO PI-R) with 240 items, the 15-item mini version is more tractable and fits into the time constraints imposed by a general household survey. Respondents were asked to indicate the degree of agreement with each statement on a 7-tier Likert-Scale from "strongly disagree" to "strongly agree." The lowest number '1' denotes a completely opposite description and the highest number '7' denotes a perfectly fitting description. Each personality trait is constructed by the average scores of three items pertaining to that trait. Thus, the value of each trait has a range between 1 to 7.

Personality traits are collected in waves 2012, 2013, 2017 and 2019 of the GSOEP. Our analysis includes individuals for whom personality traits were measured at least once. When there are multiple measurements, we average the values across the waves.[22] We standardize personality traits and use Z-scores in our empirical analysis.[23]

Cognitive ability. Cognitive skills are measured using a symbol correspondence test in the GSOEP called the SCT, which was modeled after the symbol digit-modalities-test. This test is intended to be a test of "cognitive mechanics," measuring the capacity for information processing (speed, accuracy, processing capacity, coordination and inhibition of cognitive processes).[24] Cognitive ability tests were administered in years 2012 and 2016. We include individuals for whom cognitive ability was measured at least once. When there are multiple measures, we use the average value across the waves. We standardize the cognitive ability in the same way as for personality traits and use Z-scores in the empirical analysis.

Hourly wages. The wage is calculated from self-reported gross monthly earnings and weekly working hours. Gross monthly earnings refer to wages from the principal occupation

[21] For example, a meta-analysis by Fraley and Roberts (2005) reveals a remarkably high rank-order stability: test-retest correlations (unadjusted for measurement error) are about 0.55 at age 30 and then reach a plateau of around 0.70 between ages 50 and 70.

[22] According to Roberts et al. (2008), changes of personality traits observed over a short time frame are usually attributable to noise. Therefore, we treat differences observed within a 7-year time frame to be likely from measurement error rather than fundamental changes.

[23] Z-scores are calculated by subtracting the overall sample mean (including both men and women) and dividing by the sample standard deviation. The standardized variable has mean 0 and standard deviation 1. This makes it easy to compare magnitudes of estimated model coefficients corresponding to different traits. The coefficients can also be easily interpreted as the effect of a one standard deviation change in the trait.

[24] The test was implemented asking respondents to match as many numbers and symbols as possible within 90 seconds according to a given correspondence list which is visible to the respondents on a screen. Another available test in GSOEP is a word fluency test modeled after the animal-naming-task Lindenberger and Baltes (1995): Respondents name as many different animals as possible within 90 seconds. Compared with the symbol correspondence test, this test requires sufficient language skills and therefore could be less accurate for non-native individuals. Therefore, we use SCT as our primary measure for cognitive ability.

including overtime remuneration but not including bonuses. Weekly working hours measures a worker's actual working hours in an average week.[25] The hourly wage is calculated by

$$\text{Hourly wage} = \frac{\text{Monthly gross wages (including overtime pay; without annual bonus)}}{\text{Weekly working hours} * 4.33}$$

We deflate wages using the consumer price index with 2005 serving as the base year.

Job spells and unemployment spells. Each panel wave contains retrospective monthly information about the individual's employment history. The GSOEP distinguishes between several categories of employment status, and we aggregate the information into three distinct categories: unemployed, employed and out of labor force. A person is defined as unemployed (a job searcher) if they are currently not employed and indicate that they are looking for a job. Employment status refers to any kind of working activity: full time, part time, short working hours or mini jobs. Out of labor force includes retirement, parental leave, school, vocational training and military service. We drop the spells that are out of labor force.[26] If a job A directly follows a job B in the same employment spell, we code such occurrence as a job-to-job transition. If an individual reports any unemployment spells between two jobs, then the previous job ended in a transition to unemployment. We drop individuals with missing information on key variables (education, age, gender, personality traits, cognitive ability). We further drop individuals who are out of labor force during the entire survey. The final sample contains data on 6,683 individuals.[27]

The upper panel of Table 1 presents summary statistics for labor market outcomes by gender. As seen in the last column, all of the gender differences are statistically significant at conventional levels except for the accumulated unemployment months due to its smaller sample size. Men spend fewer months in unemployment, 22.73 on average in comparison to 24.39 for women. They also spend more time in employment, 46.99 months compared to 40.66 for women. Men have on average 17.11 years full-time experience, compared to 10.24 for women. However, women have more part-time experience. The dataset contains information on actual wages. Men's actual wage is higher, €18.14 on average for men in comparison to €14.61 on average for women.

As seen in the lower panel of Table 1, the statistically significant gender wage gap occurs despite the fact that men and women have very similar average years of education (12.40 for men and 12.54 for women) and cognitive ability (3.31 for men and 3.27 for women). In terms

[25] When the actual working hours are not available, we use reported contracted working hours if available.

[26] We keep the later spells in the sample if individuals return to the labor force (or into unemployment) after having left the labor market.

[27] Appendix section A.2.1 discusses the sample selection criteria in greater detail. Table A1 compares the raw sample and the final sample.

Table 1: Summary Statistics by Gender†

	Male			Female			Difference	
	Mean	Std. Dev.	Obs.	Mean	Std. Dev.	Obs.	Diff in mean	P-value
Labor market measures								
Acc. unemployment (months)	22.73	23.48	336	24.39	23.56	428	-1.66	0.34
Acc. employment (months)	46.99	23.29	2467	40.66	24.95	2271	6.33	0.00
Full time experience (years)	17.11	11.06	4922	10.24	9.54	5315	6.88	0.00
Part time experience (years)	1.37	3.08	4922	5.83	6.62	5315	-4.45	0.00
Actual wage (€/h)	18.14	9.63	4329	14.61	8.15	4594	3.53	0.00
Demographic characteristics								
Age	43.92	9.93	2803	44.07	10.22	2699	-0.15	0.59
Cohort 1:age $\in [25, 37)$	0.25	0.43	2803	0.27	0.44	2699	-0.01	0.24
Cohort 2:age $\in [37, 49)$	0.39	0.49	2803	0.36	0.48	2699	0.04	0.01
Cohort 3:age $\in [49, 60]$	0.35	0.48	2803	0.38	0.48	2699	-0.02	0.09
Years of Education	12.40	2.82	2803	12.54	2.75	2699	-0.14	0.06
Marriage	0.68	0.47	2803	0.59	0.49	2699	0.09	0.00
Dependent child (under age 18)	1.01	1.18	2803	0.82	1.05	2699	0.19	0.00
Cognitive ability	3.31	0.91	2804	3.27	0.86	2699	0.04	0.06
Openness to experience	4.52	1.05	2804	4.72	1.06	2699	-0.20	0.00
Conscientiousness	5.79	0.78	2804	5.95	0.74	2699	-0.15	0.00
Extraversion	4.84	1.02	2804	5.10	0.98	2699	-0.26	0.00
Agreeableness	5.23	0.83	2804	5.51	0.82	2699	-0.28	0.00
Emotional stability	4.59	1.03	2804	4.12	1.10	2699	0.47	0.00

†The p-value is for a two-sided t-test of equality of means. Cognitive ability and personality traits reported in the table are their raw values. We use their Z-scores in the later empirical analysis. Observations in the upper panel reports the number of spells and Observations in the lower panel are number of individuals. each individual may have multiple spells.

of demographic characteristics, the women and men in our sample are roughly the same age on average (44). Men are more likely to be married (68 percent versus 59 percent) and to have more dependent children under the age of 18 (1.01 for men in comparison to 0.82 for women).

With regard to the Big Five personality traits, there are significant gender differences in each of the traits.[28] Women have a higher average score for all the traits except for emotional stability. Women have a lower emotional stability score by 0.47, which is the largest gender disparity observed for any of the traits. As previously noted in section one, similar personality trait differences by gender have been documented for many countries.

Comparing average hourly wages for men and women, there is a 19.5 percent gender wage gap, which is substantial considering that men and women have nearly the same years of education and cognitive skill levels. In comparison to other estimates in the literature, a

[28]In Table 1, the traits are measured on a scale of 1 to 7, as reported in the raw data. However, in our empirical analysis we use standardized z-scores for ease of interpreting effect sizes.

study by Blau and Kahn (2000) found a gender hourly gap in West Germany of 32 percent. The gap we find is consistent with reports from the German Federal Statistical Office, which show that the gender wage gap was fairly stable from 2013 to 2019, declining only slightly. The gap stood at 22 percent in 2014 and 19 percent in 2019, placing Germany as the European Union country with the second-worst gender pay gap (after Estonia).

3.1 How are personality traits associated with wages and unemployment spells

Table 2 reports associations between personality traits and log hourly wages by gender. Columns (1) and (4) report coefficients from a standard Mincer regression of log wages on education, labor market experience and its square. Column (2) and (5) reports estimated coefficients from a log wage regression that adds cognitive ability and personality traits as additional covariates. Column (3) and (6) include, in addition, marital status and dependent child status. A comparison of the coefficients from a Mincer model with ones obtained from a model that includes personality traits (e.g. columns 1 and 2 and columns 4 and 5) shows that the standard Mincer specification over-estimates the returns to education, especially for men. In specifications 3 and 6, the estimated returns to experience are very similar by gender. With regard to personality traits, three traits are correlated with hourly wages: openness to experience, agreeableness and emotional stability. Individuals with high scores on openness to experience and agreeableness tend to have lower hourly ages, whereas individuals with high scores on emotional stability receive higher hourly wages. Individuals with higher cognitive test scores also have significantly higher wages, particularly for men. The magnitudes of the statistically significant personality trait coefficients do not change much depending on whether marital and child status are included. Our estimated job search model includes cognitive scores and personality traits but not marital and child status, because these two characteristics are less direct determinants of earnings capacity.

Figure 2 displays estimated Kaplan-Meier survival functions for unemployment duration by gender. Women exit unemployment more slowly andt exit employment more quickly than men. We also have estimated a Cox proportional hazards model, shown in Table 3, to analyze how employment transitions relate to observed individual traits. For both men and women, education promotes job stability; it increases the rate of exiting unemployment and decreases the rate of exiting employment. Similarly, labor market experience reduces the hazard rate out of employment and, for women, increases the hazard rate out of unemployment. For both men and women, a ceteris paribus increase in cognitive ability increases the hazard rate out of unemployment and decreases the hazard rate out of employment.

Table 2: The association between individual traits and hourly wages (by gender)†

Outcome variable:		Male			Female	
(log) hourly wage	(1)	(2)	(3)	(4)	(5)	(6)
Years of education	0.082***	0.075***	0.075***	0.083***	0.080***	0.080***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Experience	0.047***	0.049***	0.039***	0.040***	0.040***	0.039***
	(0.003)	(0.003)	(0.004)	(0.003)	(0.003)	(0.003)
Experience^2	-0.087***	-0.086***	-0.061***	-0.070***	-0.068***	-0.062***
	(0.008)	(0.008)	(0.009)	(0.008)	(0.008)	(0.009)
Openness to experience		-0.021*	-0.013		-0.020**	-0.017*
		(0.011)	(0.011)		(0.010)	(0.010)
Conscientiousness		-0.001	-0.004		0.011	0.009
		(0.010)	(0.010)		(0.011)	(0.011)
Extraversion		0.002	0.000		0.009	0.008
		(0.010)	(0.010)		(0.011)	(0.011)
Agreeableness		-0.022**	-0.021**		-0.027***	-0.027***
		(0.010)	(0.010)		(0.009)	(0.009)
Emotional stability		0.034***	0.037***		0.024**	0.025**
		(0.010)	(0.010)		(0.010)	(0.010)
Cognitive Ability		0.070***	0.067***		0.040***	0.040***
		(0.010)	(0.010)		(0.010)	(0.010)
Constant	1.240***	1.272***	1.336***	1.111***	1.148***	1.146***
	(0.049)	(0.049)	(0.056)	(0.049)	(0.050)	(0.053)
Additional covariates						
Marriage indicator			X			X
Number of dependent children			X			X
Number of Obs	4329	4329	4329	4594	4594	4594
R^2	0.237	0.256	0.275	0.219	0.228	0.236

†Estimates based on OLS regression results. The standard errors are clustered at the individual level. Standard errors in parentheses. $p < 0.1^{*}$, $p < 0.05^{**}$, $p < 0.01^{***}$.

Figure 2: Kaplan-Meier survival estimates by gender

(a) Unemployment duration

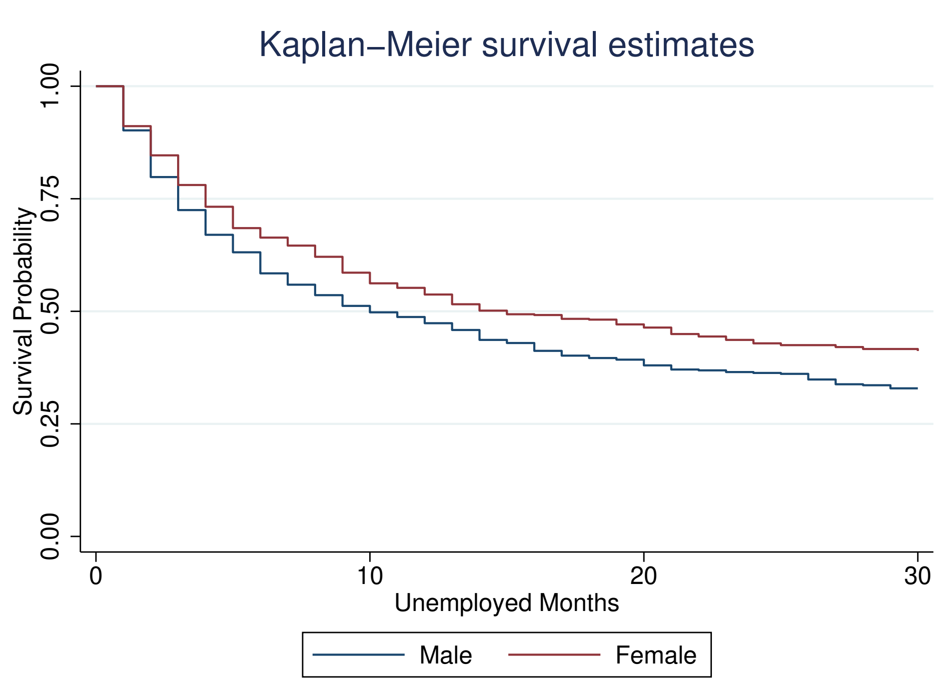

(b) Employment duration

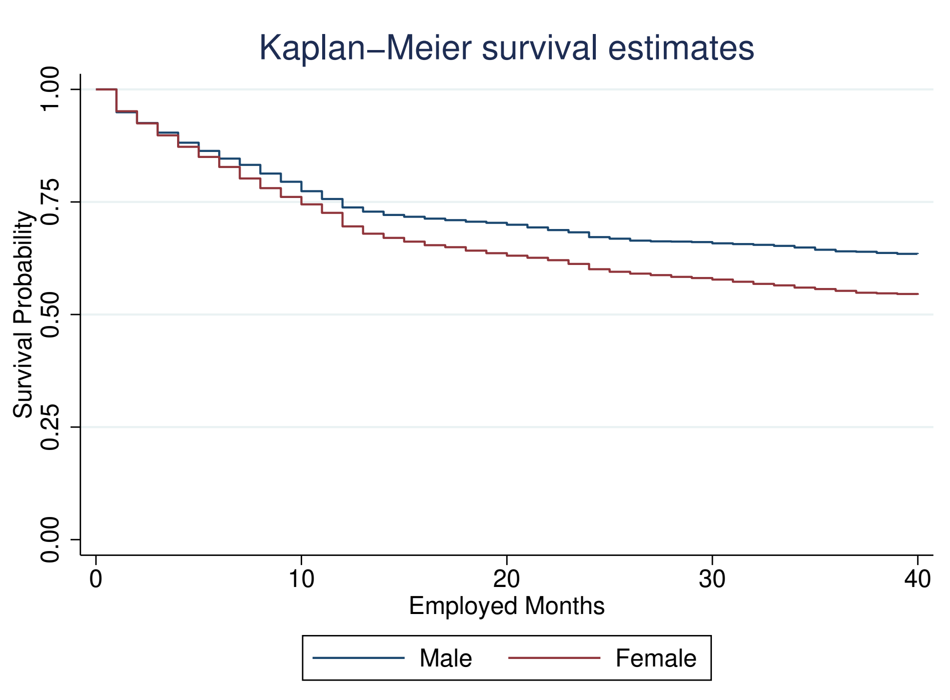

Note: Source: GSOEP data.

Table 3: Estimated unemployment and employment Cox proportional hazard rates†

Outcome variable:	Unemployment		Employment	
	(1) Male	(2) Female	(4) Male	(5) Female
Years of education	0.094***	0.227***	-0.062***	-0.017*
	(0.028)	(0.020)	(0.013)	(0.010)
Experience	0.002	0.065***	-0.113***	-0.120***
	(0.020)	(0.023)	(0.011)	(0.010)
Experience^2	0.011	-0.164**	0.156***	0.190***
	(0.055)	(0.077)	(0.029)	(0.027)
Openness to experience	0.075	-0.033	0.133***	0.065**
	(0.065)	(0.066)	(0.038)	(0.029)
Conscientiousness	0.084	0.127**	-0.097***	-0.056*
	(0.061)	(0.064)	(0.033)	(0.032)
Extraversion	-0.100	-0.053	0.093***	0.078***
	(0.068)	(0.069)	(0.035)	(0.030)
Agreeableness	-0.083	-0.074	0.016	-0.022
	(0.058)	(0.061)	(0.036)	(0.030)
Emotional stability	0.068	0.171***	-0.120***	-0.038
	(0.067)	(0.066)	(0.037)	(0.028)
Cognitive Ability	0.117**	0.259***	-0.070**	0.060**
	(0.056)	(0.063)	(0.035)	(0.030)
Number of Obs	547	634	4375	4681

†Standard errors in parentheses. $p < 0.1^{*}$, $p < 0.05^{**}$, $p < 0.01^{***}$.

Four of the personality traits (all except agreeableness) are associated with labor market transition rates. For both men and women, openness to experience and extraversion increase the hazard rate out of employment, whereas conscientiousness and emotional stability decrease it. For women, conscientiousness and emotional stability are associated with an increased hazard rate out of unemployment.

To summarize, the hourly wage regression results and the hazard model analysis demonstrate that both cognitive and noncognitive traits are important determinants of hourly wages and employment transitions. Also, focusing only on cognitive traits, such as years of education and cognitive test scores, without considering personality traits can lead to misleading inferences about the sources of wage gaps. This empirical analysis is not informative, however, about the mechanisms through which individual traits affect wages, employment and job and unemployment spell durations. For a richer understanding of how personality traits affect labor market outcomes and their importance in accounting for gender wage gaps, we next turn to the estimation of the job search model presented in Section 2.

4 Identification and estimation

In this section, we discuss the empirical implementation of the job search model. We first describe how and why we introduce measurement error in wages into the model. Next, we consider the identification of the model parameters. Lastly, we describe our maximum likelihood estimation approach.

4.1 Measurement error

We introduce measurement error in wages in the job search model for multiple reasons. First, survey data on wages typically include measurement error. In a well-known validation study using data from the Panel Study of Income Dynamics (PSID), Bound et al. (1994) find that measurement error is not a major concern in self-reported annual earnings measures. However, they find that reported hourly wage compensation contains a greater degree of measurement error, with the proportion of log wage variation attributable to measurement error reaching 50 to 60 percent. The GSOEP respondents probably report their monthly earnings more accurately than the PSID respondents, as they are required to have their pay statements on hand at the time of reporting their salaries. However, reports on hours worked may be subject to a greater degree of measurement error. Also, rounding errors, recall bias and social desirability bias may all contribute to measurement error in survey data.

A second reason for incorporating measurement error into the model is to ensure that our likelihood is not degenerate. In the data, the vast majority of (direct) job-to-job transitions are associated with wage increases as implied by the theoretical model. However, there are also some reported wage decreases associated with job-to-job movements. With measurement error, the likelihood of observing a wage decrease is positive.[29] Additionally, the theoretical model yields a single wage within each job spell, whereas in the data we may observe some wage variation within spells, which can also be explained by measurement error.

In addition to accounting for the possibility of decreasing wages happening both within and between job spells, we also allow for measurement error in wages for individuals transiting from unemployment to a job paying a wage w. The reservation wage of an individual depends

[29]There are some alternative theoretical models studied in the literature that can generate wage decreases. Two such examples are Postel-Vinay and Robin (2002) and Dey and Flinn (2005). In Postel-Vinay and Robin, workers may take a wage reduction to move to a "better" firm because of the increased future bargaining advantage being at that firm conveys. In Dey and Flinn, in addition to wages, firms and workers profit from the worker having health insurance. When a worker moves from a firm in which they do not have health insurance to one in which they do, their bargained wage may decrease since they share the cost of the health insurance with the firm in proportion to their share,α, of the job surplus. However, models based on the assumption of Bertrand competition between firms generate the implication that real wages over a job spell at an employer should never decrease, which is observed in the data.

on all of the primitive parameters of the model, and for each individual i we are allowing their parameters to be a function of z_i and the estimated values of the γ vector associated with each of the primitive parameters. The accepted wage for an individual with characteristics z_i may be less than the value of their reservation wage implied by our model. Measurement error is required for this to be a positive probability event.

As is commonly done, we assume a classical measurement error structure. That is, we assume that observed wages $\tilde{w}$ are

$$\tilde{w} = w\varepsilon$$

where $\tilde{w}$ is the reported wage and w is the worker's "true" wage. We assume that the measurement error, ε, is independently and identically distributed both within individuals across job spells and across individuals and that it is lognormal (Wolpin (1987); Flinn (2002)). The density of ε is

$$m(\varepsilon) = \phi\left(\frac{\log(\varepsilon) - \mu_\varepsilon}{\sigma_\varepsilon}\right) / (\varepsilon\sigma_\varepsilon) \tag{11}$$

where ϕ denotes the standard normal density and μ_ε and σ_ε are the mean and standard deviation of $log(\varepsilon)$. We impose the restriction $\mu_\varepsilon = -0.5\sigma_\varepsilon^2$, so that $E(\varepsilon|w) = 1$.[30] The expectation of the observed wage is equal to the true wage, because

$$E(\tilde{w}|w) = w \times E(\varepsilon|w) = w \;\forall w.$$

The measurement error dispersion parameter, σ_ϵ, can be identified from multiple wage measures within the same job spell. Let $\tilde{w}_k^t$ and $\tilde{w}_k^{t+1}$ denote wage measures at two different periods in the same job spell k. Our model implies that the difference in the two values is produced solely by measurement error, because the model implies that the actual wage is invariant within a job spell. Given the multiplicative structure, $\tilde{w} = w\epsilon$, we have

$$\log \tilde{w}_k^{t+1} - \log \tilde{w}_k^t = \log \epsilon^{t+1} - \log \epsilon^t \sim N(0, 2\sigma_\epsilon^2)$$

We can therefore identify σ_ϵ from the distribution of observed wage variation within job spells. We plot the empirical distribution of $\log \tilde{w}_k^{t+1} - \log \tilde{w}_k^t$ in figure 3. The distribution shows an average growth rate close to 0, with significant dispersion.[31]

In constructing the likelihood, when there are multiple wage measures within a job spell,

[30] Given ϵ follows a lognormal distribution, $E(\epsilon) = \exp\left(\mu_\epsilon + 0.5\sigma_\epsilon^2\right) = 1$ if $\mu_\epsilon = -0.5\sigma_\epsilon^2$.

[31] This pattern is more supportive of the no-renegotiation model than the renegotiation model (described in section 2.2). With the renegotiation model, we would expect to observe a weakly positive wage growth rate when incumbent firms increase a worker's wage in order to retain them.

Figure 3: The distribution of annual real wage growth

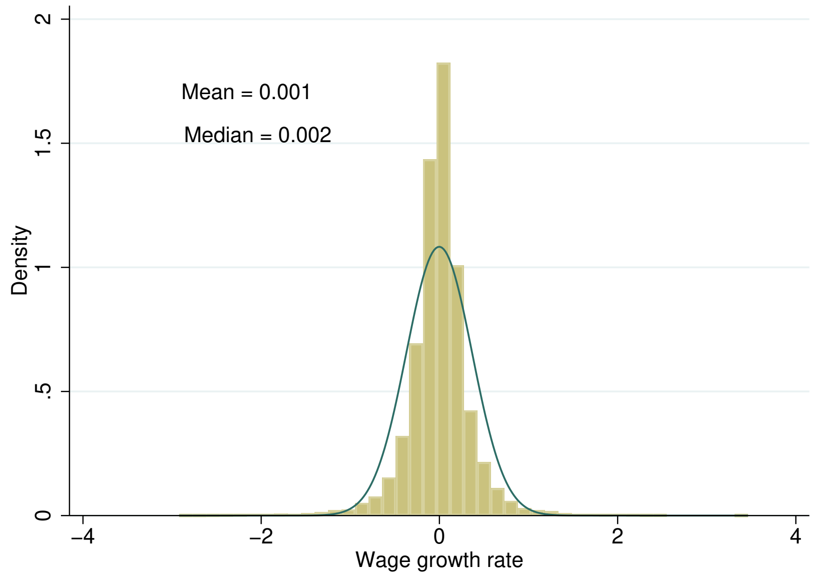

we use the average wage, taking into account the number of observations used to construct the average. Thus, the measurement error associated with the average ln wage in the k^{th} job spell is distributed as

$$\bar{\epsilon}_k \sim N(0, \frac{\sigma_\epsilon^2}{n}), \tag{12}$$

where n is the number of ln wage observations for the spell.

4.2 Identification

We next provide a brief discussion of how the model parameters, including ability, a, the bargaining parameter, α, the match quality, θ, and the transition parameters $\{\lambda_E, \lambda_U, \eta\}$ are separately identified. Appendix A.3 considers identification in greater detail, building on earlier identification arguments in Flinn and Heckman (1982) and Flinn (2006).

The analysis in Flinn and Heckman (1982) considers the estimation of a nonequilibrium search model with an exogenous wage offer distribution, which can be thought of as a special case of the model developed in this paper when $\alpha = 1$. They consider the homogeneous case in which the primitive parameter values are equal for all labor market participants. Furthermore, they assume that wages are measured without error, and that there is no on-the-job search. They demonstrate that the parameters λ_U, η, and the parameters characterizing the population wage offer distribution are identified only using data from the basic monthly Current Population Survey. These data are limited to the wages for currently employed individuals and the length of on-going unemployment spells for those who are unemployed at the survey date. They further show that the flow utility of unemployment b and the instantaneous discount rate ρ are not point identified. Assuming a value of either one of them, however, enables point identification of the other.

The assumptions that there is no measurement error in wages, that the labor market environment is stationary, and that all individuals share the same labor market parameters, allow Flinn and Heckman to concentrate out of the likelihood function the reservation wage w^* by replacing it with its (super) consistent estimator $\widehat{w^*} = \min[w_{(1)}, w_{(2)}, ..., w_{(N_E)}]$, where N_E is the number of employed people with wage observations. In this case, all accepted wages must be no greater than the (common) reservation wage estimate for the sample. This estimation method is not possible to use in this case, because we have continuously varying reservation wages in the population.

The identification of the surplus share parameter α was considered in some detail in Flinn (2006). In a homogeneous stationary model with no on-the-job search but with bargaining, a sufficient condition for the identification of the surplus share parameter α is that the distribution $G(\theta)$ not belong to parametric location-scale family.[32] Under the assumption of lognormality, the wage distribution is not location-scale, although $\log w$ is. This relatively slight nonlinearity is enough to enable identification of α in theory. In practice, relatively large samples are required to precisely estimate α.

Adding on-the-job search to the model is not at all problematic, in principle, if the data are longitudinal and job-to-job transitions can be observed. The rate of leaving one job paying a wage of w for another paying w', where $w' > w$, is simply $\lambda_E(1 - G(w))$. Of course, these w and w' are the "true" wages, and we need to account for measurement error when formulating the likelihood function. This is straightforward under the classical measurement error assumption, as explained below.

Our analysis differs from the previously cited ones by the inclusion of an individual ability parameter a. We now turn to the identification and estimation of this parameter. Equation 8 specifies how wages are determined by ability (a), bargaining power (α), match quality (θ) and the reservation match quality (θ^*):

$$w(\theta; a) = a\left(\alpha\theta + (1-\alpha)\theta^*\right), \theta \geq \theta^*$$

[32]In the analysis of Flinn, individual variability in productivity, a, was ignored. This is due to the fact that the CPS data used were cross-sectional, making the identification of the separate distributions of a and θ impossible except under stringent parametric assumptions on both and access to extremely large samples. We consider the estimation of the distribution of a in detail below

Both ability a and bargaining power α increase wages. The transition parameters $\{\lambda_E, \lambda_U, \eta\}$ and the parameters describing the distribution of θ affect wages only indirectly through their impact on θ^*.

As discussed in Section 2, the value of θ^* is independent of a, so that we have

$$\log w(\theta; a) = \log a + \log\left(\alpha\theta + (1-\alpha)\theta^*\right), \theta \geq \theta^* \tag{13}$$

The log wage consists of an additive individual fixed effect, $\log a$, and a term that involves the current match quality value, θ, the individual's surplus share parameter, α, and the reservation match quality θ^*.[33]

Recall that in the heterogeneous model specification, each of the primitive parameters depends on an index function $z_i\gamma_j$, where j is the subscript denoting the relevant primitive parameter. The reservation match value for individual i is given by $\theta^*(z_i, \gamma_{-a})$, where γ_{-a} denotes the matrix of index function parameters excluding γ_a. Recall that $a_i = \exp(z_i\gamma_a)$, so that

$$\log \tilde{w}_i = z_i\gamma_a + \Omega(\theta; z_i, \gamma_{-a}) + v_i, \tag{14}$$

where $\Omega(\theta; z_i, \gamma_{-a}) = \ln(\text{logit}(z_i\gamma_\alpha) \times \theta + (1 - \text{logit}(z_i\gamma_\alpha)) \times \theta^*(z_i, \gamma_{-a}))$ and where $v_i \equiv \ln \varepsilon_i + 0.5\sigma_\varepsilon^2$ is distributed as $N(0, \sigma_\varepsilon^2)$ since $\ln \varepsilon_i$ is distributed as $N(-0.5\sigma_\varepsilon^2, \sigma_\varepsilon^2)$. This means that the intercept term in the ln wage equation is equal to the intercept term $\gamma_{a,0}$ in the index function associated with a plus the constant $0.5\sigma_\varepsilon^2$.[34]

For individuals whose job spell followed an unemployment spell, the density of θ is a truncated distribution given by

$$\frac{g(\theta; z_i)}{\tilde{G}(\theta^*(z_i, \gamma_{-a}); z_i)}, \ \theta \geq \theta^*(z_i, \gamma_{-a}).$$

In this case, the conditional expectation of $\ln \tilde{w}_i$ is

$$E(\ln \tilde{w}_i | z_i) = z_i\gamma_a + \int \Omega(\theta; z_i, \gamma_{-a}) \times \frac{g(\theta; z_i)}{\tilde{G}(\theta^*(z_i, \gamma_{-a}); z_i)} d\theta.$$

We see that the first term $z_i\gamma_a$ does not include any other index function parameters, while

[33] Postel-Vinay and Robin (2002) perform a similar log-linear wage decomposition.

[34] Since the variance of the measurement error σ_ε^2 can be consistently (and unbiasedly) estimated from the differences in reported ln wages on the same job, a consistent estimator of $\gamma_{a,0}$ can be obtained from the estimated intercept term in the ln wage equation.

the second term is a nonlinear function of all of the index function parameters with the exception of γ_a. Given the nonlinearities implied by the model and functional form assumptions, the parameters comprising γ_a and γ_{-a} could be in principle estimated using nonlinear least squares.

There is additional identifying information that comes from the fact that not all measured wages in the data are associated with first jobs in an employment spell. For each individual i, there exist marginal distributions of θ associated with each job within an employment spell. Let $\theta_{(j)}$ denote the value of θ associated with j^{th} job in an employment spell. The $\theta(j)$ distributions will exhibit stochastic dominance in the sense that $G_{(j)}(\cdot; z_i) \prec_{FOSD} G_{(j+1)}(\cdot; z_i)$ for $j = 1, 2, \ldots$. The conditional expectation of $\ln \tilde{w}_i$ will be a function of the γ as well as the order of the job in the employment spell. The conditional mean wage equations associated with different order jobs within an employment spell provide additional moment equations for identifying model parameters γ_a and γ_{-a}.

In moderately sized samples such as ours, it would be difficult to precisely estimate γ_a and γ_{-a} using only wage data. However, our data also contain information on labor market transitions, which, as shown above, depend on the subset of the model parameters, γ_{-a}. The labor market transition parameters can be identified from their empirical counterparts in the data (i.e. we observe the transition probabilities for individuals with different observables z). Using both wage and labor market transition data, we can obtain precise estimates of γ_a and γ_{-a} even with moderately-sized samples.[35]

As described below, the estimation approach that we adopt is maximum likelihood. The likelihood efficiently uses all of the available information on wages and labor market transition dynamics, and this is required in order to precisely estimate the relatively large number of parameters in our model. A detailed discussion of the identification requirements using maximum likelihood can be found in Appendix A.3. A key requirement is the usual full column rank condition on the matrix Z which ensures the the Hessian associated with the ln likeihood function will be of full rank. We have not discussed in detail identification under the particular function form assumptions that are made in linking the linear index functions with the primitive parameters; more detail on this topic is contained in A.3. The maximum likelihood estimates of the parameters are derived from the usual score functions. Although we have not shown that the ln likelihood function is globally concave, our attempts to investigate how varying starting values would impact the estimated parameters always

[35] One potential estimator would use the transition data to identify all of the parameters in γ_{-a}, and then to estimate the $\ln \tilde{w}$ linear regression after substituting the estimates of γ_{-a} into the expression for the expected value of θ given the individual's value of z_i and the order of the job spell in the employment spell. While appealing, the efficiency of this estimator will be dominated by our full-information maximum likelihood estimator, at least in large samples.

led to the same set of estimates.

4.3 Constructing the individual likelihood contribution

We estimate the model parameters using maximum likelihood. In this subsection, we first discuss how we construct each individual likelihood $L_i, i = \{1, 2, .., N\}$ conditional on an individual-specific set of primitive parameter values Ω_i. In the next subsection, we describe the mapping between individual characteristics z_i and Ω_i. For notational simplicity, we suppress the individual subscript i, but the reader should bear in mind that the underlying econometric model allows the search-environment parameters to vary across individuals.

As in Flinn (2002) and Dey and Flinn (2005), for example, the information used to construct the likelihood function is defined as an employment cycle. The exact composition of an employment cycle depends on the individual's initial status. If an individual enters into our sample with an existing job, the first employment cycle begins with this existing job, followed potentially by more jobs, and ends by a transition into the unemployment status. If an individual is initially unemployed, then the employment cycle begins with an unemployment spell, followed by one or more jobs, and ended by a transition into unemployment. For computational simplicity, we limit our attention to the first two jobs in the employment spell. That is, an employment cycle (EC) can consist of

$$EC = \begin{cases} \underbrace{\{t_k, \tilde{w}_k, q_k, r_k\}_{k=1}^2}_{\text{Up to two consecutive jobs}} & \text{One employment spell with a pre-existing job} \\ \underbrace{\{t_U, r_U\}}_{\text{Unemployment spell}}, \underbrace{\{t_k, \tilde{w}_k, q_k, r_k\}_{k=1}^2}_{\text{Up to two consecutive jobs}} & \text{One unemployment spell + one employment spell} \end{cases}$$

Each individual may contribute information on multiple employment spells to the likelihood. Depending on the initial status upon entering into the sample, the first employment cycle may consist of the current (incomplete) employment spell or an (incomplete) unemployment spell followed by an employment spell. The following employment cycle would always be a combination of one completed unemployment spell and one completed employment spell. For the unemployment state, t_U is the length of the unemployment spell and r_U is an indicator variable that takes the value 1 if the unemployment spell is right-censored. In the following employment spell, which consists of up to 2 jobs, for each job spell $k \in \{1, 2\}$, t_k is the length of job k in the employment spell, $\tilde{w}_k$ is the observed wage in job k, and $r_k = 1$ indicates that the duration of job k is right-censored. Lastly, q_k indicates whether the $k-th$ job ends with a transition into unemployment ($q_k = 0$) or a transition into the next job ($q_k = 1$).

We now discuss the question of left-censoring in the data, which arises because the first unemployment/employment spells began before the sample period. In the case of unemployment, we only observe the portion of the unemployment spell that began when the observation period did. Due to the stationarity assumption underlying the model, for initially unemployed workers this left-censoring can be ignored due to the "memoryless" property of the exponential distribution; that is, if job offer arrivals follow a Poisson process, the distribution of the remainder of time in the unemployment state (i.e., the forward recurrence time) is independent of the elapsed time spent in unemployed search prior to the onset of the sample period. In the case of sample members who were employed at the beginning of the sample period, we only have access to a (noisy) observation of their wage at the beginning of the sample period, and have no other information regarding their previous jobs held during the sampled employment spell. Given the structure of the model, the workers initial "true" wage rate at the time of the sample, w_1, is a sufficient statistic for their labor market outcomes following the survey date. In other words, conditional on wage rate w_1, the job history during the current employment spell before the sample period began has no impact on labor market outcomes during the remainder of the employment spell.

In describing the individual likelihood contribution, it is useful to distinguish between two types of employment cycles. An employment cycle starting with an unemployment spell can be one of the following six cases:

1. One right-censored unemployment spell ($r_U = 1$)
2. One completed unemployment spell ($r_U = 0$)
   (a) + first right-censored job spell ($r_1 = 1$)
   (b) + first completed job spell ending with unemployment ($r_1 = 0, q_1 = 0$)
3. One completed unemployment spell + first completed job spell ($r_1 = 0, q_1 = 1$)
   (a) + second right-censored job spell ($r_2 = 1$)
   (b) + second completed job spell ending with unemployment ($r_2 = 0, q_2 = 0$)
   (c) + second completed job spell ending with third job ($r_2 = 0, q_2 = 1$)

The following likelihood function expression for the employment cycle includes all six cases:

$$\begin{aligned}
l(t_U, r_U, \tilde{w}_1, t_1, r_1, q_1, \tilde{w}_2, t_2, r_2, q_2) &= \textstyle\int_{w^*} \int_{w_1} h_U^{(1-r_U)} \exp(-h_U t_U) \\
&\times \left\{ \exp\left(-h_E(w_1)t_1\right) \left[\left(\lambda_E \bar{F}(w_1)\right)^{1-q_1} \eta^{q_1} \right]^{1-r_1} \tfrac{1}{w_1} m\left(\tfrac{\tilde{w}_1}{w_1}\right) \right\}^{1-r_U} \\
&\times \left\{ \exp\left(-h_E(w_2)t_2\right) \left[\left(\lambda_E \bar{F}(w_2)\right)^{1-q_2} \eta^{q_2} \right]^{1-r_2} \tfrac{1}{w_2} m\left(\tfrac{\tilde{w}_2}{w_2}\right) \right\}^{1-r_1} \\
&\times \tfrac{f(w_1)}{\bar{F}(w^*)} \tfrac{f(w_2)}{\bar{F}(w_1)} dw_2 dw_1
\end{aligned} \tag{15}$$

where the reservation wage, w^*, is given in equation (4). The measurement error density $m(.)$ is defined in equation 12. The wage distribution $F(w)$ is based on the one-to-one mapping from the distribution of match quality $G(\theta)$ derived in equation 9, and $\tilde{F}(w) \equiv 1 - F(w)$. Additionally, h_U and $h_E(w)$ are the hazard rates when workers are unemployed and employed, respectively. They are defined as:

$$\begin{aligned} h_U &= \lambda_U \tilde{F}(w^*) = \lambda_U \tilde{G}(\theta^*) \\ h_E(w) &= \eta + \lambda_E \tilde{F}(w), \end{aligned}$$

We compute the likelihood function by Monte Carlo integration using importance sampling.[36]

An employment cycle starting with a pre-existing employment spell can be one of the following five cases:

1. One right-censored job spell ($r_1 = 1$)
2. One completed job spell ending with unemployment ($r_1 = 0, q_1 = 0$)
3. One completed job spell ($r_1 = 0, q_1 = 1$)
   (a) + second right-censored job spell ($r_2 = 1$)
   (b) + second completed job spell ending with unemployment ($r_2 = 0, q_2 = 0$)
   (c) + second completed job spell ending with third job ($r_2 = 0, q_2 = 1$)

The following likelihood expression for employment cycles that begin with a pre-existing employment spell with a initial wage w_1:

$$\begin{gathered} l(\tilde{w}_1, t_1, r_1, q_1, \tilde{w}_2, t_2, r_2, q_2 | w_1) = \\ \int_{w_1} \left\{ \exp\left(-h_E(w_1)t_1\right) \left[\left(\lambda_E \bar{F}(w_1)\right)^{1-q_1} \eta^{q_1} \right]^{1-r_1} \right\}^{1-r_U} \\ \times \left\{ \exp(-h_E(w_2)t_2) \left[\left(\lambda_E \bar{F}(w_2)\right)^{1-q_2} \eta^{q_2} \right]^{1-r_2} \tfrac{1}{w_2} m\left(\tfrac{\tilde{w}_2}{w_2}\right) \right\}^{1-r_1} \\ \tfrac{f(w_2)}{\bar{F}(w_1)} dw_2 \end{gathered}$$

Because of measurement error, the initial wage w_1 is not directly observed, so the unconditional likelihood function with a pre-exiting employment spell needs to integrate over the

[36] We generate 2500 repetitions of the (w_1, w_2) draws (50 draws of w_1 and 50 draws of w_2) for use in the importance sampling algorithm.

measurement error distribution:

$$
\begin{aligned}
&l(\tilde{w}_1, t_1, r_1, q_1, \tilde{w}_2, t_2, r_2, q_2) \\
&= \int_{w^*} l(\tilde{w}_1, t_1, r_1, q_1, \tilde{w}_2, t_2, r_2, q_2 | w_1) \tfrac{\tilde{w}_1}{w_1^2} m\left(\tfrac{\tilde{w}_1}{w_1}\right) \tfrac{1}{\Gamma(\tilde{w}_1)} dw_1 \\
&= \int_{w^*} \int_{w_1} \left\{ \exp\left(-h_E(w_1)t_1\right) \left[\left(\lambda_E \bar{F}(w_1)\right)^{1-q_1} \eta^{q_1}\right]^{1-r_1} \tfrac{\tilde{w}_1}{w_1^2} m\left(\tfrac{\tilde{w}_1}{w_1}\right) \right\} \\
&\times \left\{ \exp(-h_E(w_2)t_2) \left[\left(\lambda_E \bar{F}(w_2)\right)^{1-q_2} \eta^{q_2}\right]^{1-r_2} \tfrac{1}{w_2} m\left(\tfrac{\tilde{w}_2}{w_2}\right) \right\}^{1-r_1} \\
&\tfrac{1}{\Gamma(\tilde{w}_1)} \tfrac{f(w_2)}{\bar{F}(w_1)} dw_2 dw_1
\end{aligned}
\tag{16}
$$

where $\frac{\tilde{w}_1}{w_1^2} m\left(\frac{\tilde{w}_1}{w_1}\right) \frac{1}{\Gamma(\tilde{w}_1)}, w_1 > w^*$ denotes the density of the "true" wage based on the its measured wage; the term $\frac{\tilde{w}_1}{w_1^2}$ is the Jacobian of the transformation, $m(.)$ is the density function of the measurement error defined in equation (11), and $\Gamma(\tilde{w}_1)$ is a normalizing constant which ensures that the density integrates to unity.

We then construct the overall log likelihood function L for the whole sample (of size N). Our model assumes that an individual i has their individual-specific set of labor market parameters $\Omega_i = \{\lambda_U(i), \lambda_E(i), \alpha(i), \eta(i), a(i), b(i), \sigma_\theta(i)\}$. As discussed below, these parameters are gender-specific functions of observable heterogeneity represented by a row vector of characteristics z_i, which includes education, cognitive skills, personality traits and birth cohort. The log likelihood function $\ln L$ defined for the entire sample is

$$
\ln L = \sum_{i=1}^{N} \sum_{j=1}^{J} \ln l_{ij}(\text{Employment cycle}_{ij} | \Omega_i)
$$

where $l_{ij}(\text{Employment cycle}_{ij} | \Omega_i)$ is the likelihood function for the $emphitj^th$ employment cycle for individual i defined by either equation (15) or (16). Because individual heterogeneity is (essentially) continuously distributed, computing individual i's log likelihood contribution at each iteration of the estimation algorithm requires solving for each person's reservation wage strategy separately.

5 Model estimates

5.1 Estimated model parameters under homogeneous/heterogeneity specifications

Many previous papers have estimated search models that allow parameters to differ by gender (e.g. Bowlus (1997), Bowlus and Grogan (2008), Flabbi (2010a), Liu (2016), Morchio

Table 4: Parameter estimates under alternative heterogeneity specifications

Parameter	Description	Homogeneous		Heterogeneous†	
		Male	Female	Male	Female
a	ability	20.20	16.93	19.74	18.18
		(0.26)	(0.21)	(2.23)	(1.84)
α	surplus division (bargaining)	0.481	0.438	0.481	0.406
		(0.008)	(0.008)	(0.039)	(0.056)
η	separation rate	0.003	0.002	0.003	0.003
		(0.000)	(0.000)	(0.003)	(0.002)
λ_U	offer arrival rate, in unemployment	0.145	0.143	0.160	0.162
		(0.002)	(0.001)	(0.067)	(0.074)
τ	offer arrival rate ratio (emp/unemp)	0.318	0.480	0.350	0.484
		(0.007)	(0.010)	(0.002)	(0.010)
b	flow utility when unemployed	-1.940	-1.453	-1.493	-1.419
		(0.013)	(0.021)	(0.007)	(0.011)
σ_θ	$\ln \theta \sim N\left(-\frac{\sigma_\theta^2}{2}, \sigma_\theta^2\right)$	0.446	0.424	0.426	0.404
		(0.005)	(0.005)	(0.004)	(0.004)
σ_ϵ	$\ln \epsilon \sim N\left(-\frac{\sigma_\epsilon^2}{2}, \sigma_\epsilon^2\right)$	0.178	0.175	0.149	0.141
		(0.002)	(0.002)	(0.001)	(0.001)
N		6,483		6,483	
$\log L$‡		-53,008		-51,096	
P-value (LR tests)				< 0.001	

†In the heterogeneous specification, the parameters $\{a, \alpha, \lambda_U, \eta\}$ depend on indices of individual characteristics. For these parameters, the number in parenthesis reports the standard deviation of the parameter distribution rather than standard errors.
‡The likelihood ratio (LR) test tests the current specification against the previous one. The monthly discount rate is set at 0.005.

and Moser (2020), Amano-Patino et al. (2020)). However, gender is only one of many individual traits that may be relevant to the job search process and labor market outcomes. The index formulation introduced in the previous section allows for a greater degree of individual heterogeneity, with parameters depending on gender, education level, cognitive skills, age (birth cohort), and personality traits.

Table 4 reports the estimated search model coefficients for a "homogeneous" specification, in which model parameters only differ by gender but are otherwise assumed to be homogeneous within each gender, and a "heterogeneous" specification, in which the parameters are allowed to be gender-specific and, in addition, to potentially vary by education, cognitive skills, personality traits, and age cohort. Figure 4 shows the distributions of the estimated parameter values for males and females under the heterogeneous model and Table 4 shows the mean parameter values (in the last two columns).

A comparison of the estimates for the homogeneous and heterogeneous specifications

Figure 4: The distribution of search parameters $\{a, \alpha, \lambda, \eta\}$

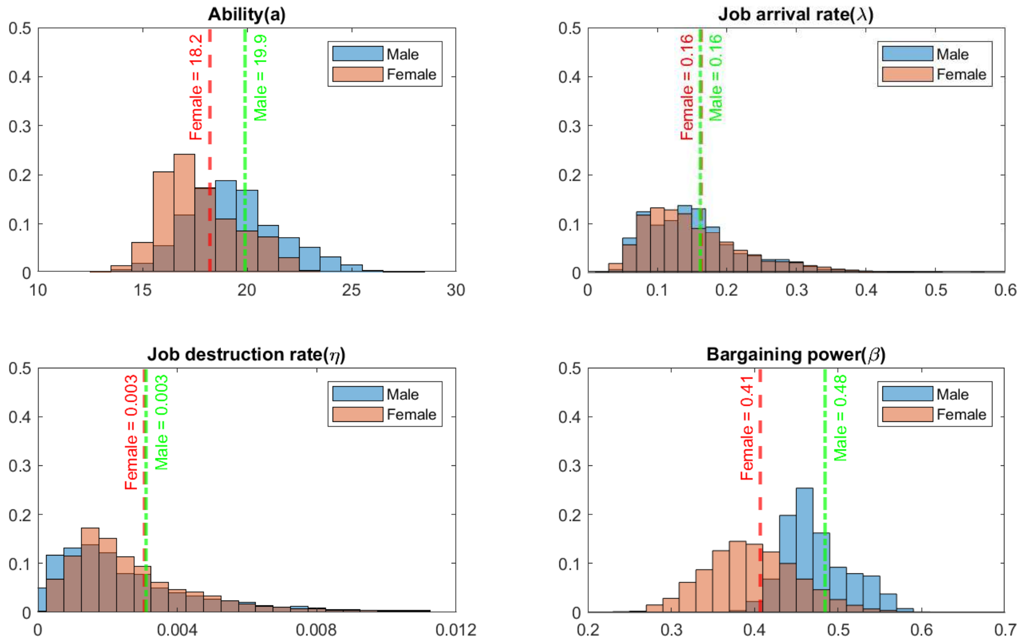

reveals important gender differences as well as substantial individual heterogeneity. First, the estimated ability parameters (a) indicate that males are more productive on average than females.[37] Average female productivity is 16.93 in comparison to 20.20 for men for the homogeneous model and 18.18 in comparison to 19.74 for the heterogeneous model. The 9 to 16 percent gap is slightly smaller than gaps found in other studies based on U.S. data. For example, Bowlus (1997) finds the productivity of females is 17 percent lower using data from the NLSY79, and Flabbi (2010a) finds a 21 percent differential using CPS data. As seen in Figure 4, there is substantial variance in the estimated ability parameters and the male and female distributions overlap.

In addition to having higher estimated average productivity values, men are estimated to have a higher surplus division parameter (α), so that they receive a larger share of the job surplus than do women, on average. The estimated parameter values range from 0.4 to 0.5, which is fairly consistent with values reported in the search literature using similar modeling frameworks.[38] For example, Bartolucci (2013) uses German matched employer-employee data

[37]Total productivity is $y = a \times \theta$. We set the location parameter of the match value distribution to be $\mu = -0.5\sigma_\theta^2$ so that $E[\theta] = 1$. Therefore, $E[y] = E[a\theta] = E[a]$. Note, however, that in accepted matches $E[\theta] > 1$ since only $\theta \geq \theta^*$ are accepted.

[38]When Bertrand competition between firms is assumed, the estimated surplus share parameter is typically significantly lower. This is due to the fact that wages are bid up as firms directly compete for workers, meaning that the average share of the surplus received by workers is an increasing function of λ_E, which is

Figure 5: Model goodness of fit to wage distributions

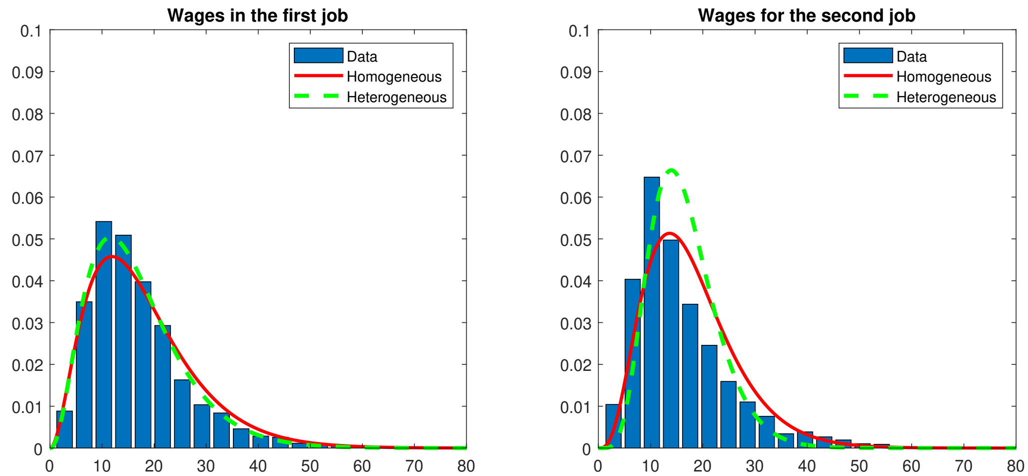

and finds female workers have, on average, slightly lower bargaining power than their male counterparts, with an average α of 0.42 (for both genders). Flinn (2006), using CPS data, finds that the overall bargaining power is approximately 0.45. Figure 4 shows substantial heterogeneity in bargaining parameters across individuals, again with substantial overlap in the male and female distributions.

Estimated job separation rates η are generally small in magnitude and are estimated to be similar for men and women (0.002-0.003). Since the time unit used is weeks, this implies that the average length of time until a job is exogenously terminated is approximately 6 years. Of course, jobs will end sooner than that, on average, since workers also leave jobs for better wages at alternative employers. Figure 4 shows that the distribution of job separation rates exhibits right skewness, with the vast majority of people having very low probabilities of job separation and a small fraction having more extreme values. Men have a slightly higher job offer arrival rate compared to women and a lower flow utility when unemployed.

We next consider the estimates of the standard deviations of measurement error and θ under the two specifications. By allowing for heterogeneous parameters, the model relies less on randomness components (match quality and measurement error) to fit the data. This pattern suggests that much of the variance that the homogeneous model attributes to randomness can actually be explained by observable characteristics.

The homogeneous specification is nested within the heterogeneous specification, allowing for a straightforward likelihood ratio test of the model restrictions. The p-values are reported in the two bottom lines of Table 4. The LR test strongly rejects the homogeneous specification in favor of the flexible heterogeneous one. In addition to performing the formal test,

a key determinant of the rate at which these competitions occur. See, e.g.,Flinn and Mullins (2015).

we also graphically examine whether the heterogeneous model provides a better goodness of fit. Figure 5 shows the model fit to the wage distributions on the first and second jobs. Both models fit the first job wage distributions reasonably well, but the heterogeneous model clearly does a better job of fitting the second job wage distributions. Figure 6 graphs the distributions of the unemployment spell length, and the spell lengths of the first and second jobs, both in the data and under the homogeneous and heterogeneous specifications.[39] The heterogeneous specification better captures the right skewness in the unemployment spell length distribution.

5.2 Understanding the role of personality traits and other individual characteristics in a job search model

We next examine how education and personality traits affect job search parameters $\{\lambda_U, \eta, \alpha, a\}$.[40] Table 5 reports the parameter estimates for the heterogeneous model that are informative about the channels through which education, cognitive skills, birth cohort, and personality traits influence wage and employment outcomes. For men and women, education increases the unemployment job offer arrival rate (λ_U) and lowers the job separation rate (η). Education also increases productive ability (a) and increases bargaining power (α). Conditional on education, cognitive ability significantly increases productive ability and increases job offer arrival rates for both men and women. Thus, education and cognitive ability enter through multiple model channels, which combine to increase wages and promote employment stability.

As seen in Table 5, personality traits are statistically significant determinants of job search parameters, and, for the most part, affect parameters of men and women in similar ways. As previously noted, conscientiousness has been emphasized in prior literature as the trait most strongly associated with superior labor market outcomes. Consistent with this view, our estimates indicate that conscientiousness increases job offer arrival rates and decreases job separation rates. It also increases productive ability for men and increases bargaining power for women. All of these effects contribute to higher wage levels and more stable employment. Emotional stability is another trait that leads to better labor market prospects. For women, emotional stability increases the job offer arrival rate, lowers the job separation rate, and substantially augments bargaining power. For men, emotional stability increases ability, increases job arrival rates, and lowers the job separation rate. The remaining three traits - openness to experience, extraversion and agreeableness - are not necessarily

[39] In generating these distributions, we incorporate any censoring that is in the data.

[40] Given the assumption that λ_E is proportional to λ_U, the effects of personality traits on λ_E are also proportional to the effect on λ_U.

Figure 6: Model goodness of model fit to spell length distributions

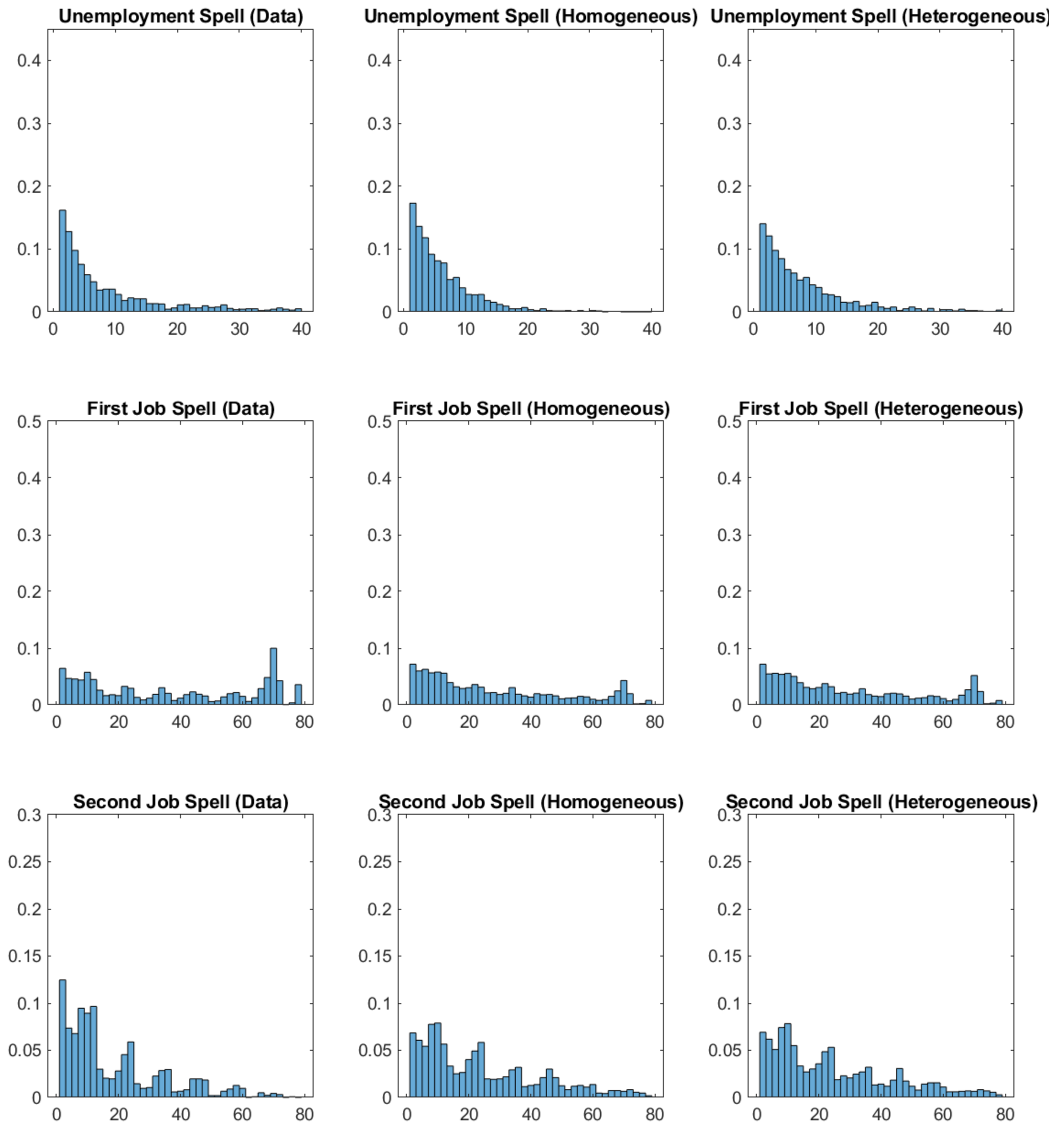

desirable characteristics from a labor market perspective. For men, openness to experience increases the job separation rate and decreases bargaining power. For women, it decreases job exit rates but also negatively affects bargaining power. The extraversion trait appears to only be important for men; it increases job offer rates, increases job separation rates, and decreases bargaining power. Lastly, agreeableness has a uniformly negative effect on labor market parameters. For both men and women, agreeableness lowers job offer arrival rates and increases job separation rates. It also lowers ability for men and lowers bargaining power for women.

The job search model that we estimate is stationary, and we therefore do not condition on time-varying state space elements (such as labor market experience). However, we include birth cohort indicators to capture possible differences in the labor markets for workers of different ages. As seen in the bottom rows of Table 5, younger workers have lower ability and much higher job separation rates than middle age workers (the reference category is age 37-48). For men, younger workers have a significantly higher job offer arrival rate but for women the job offer arrival rate is lower. Older workers (age 49-60) are estimated to have lower ability, much lower job offer arrival rates, and lower job separation rates. For women, being older significantly reduces bargaining power, while older age men do not have such a penalty.

6 Interpreting the model estimates

In this section, we use the estimated model to analyze how different cognitive and noncognitive traits affect labor market outcomes and the implications for gender disparities. The model is a steady state model, so we first simulate outcomes for 500 periods (months) to get a reasonable approximation to the steady state.

6.1 Effects of cognitive and noncognitive traits on wage and employment outcomes

Table 6 explores the effects of a ceteris paribus change in each of the individual traits on labor market outcomes. Specifically, we increase each trait by one standard deviation for all individuals (holding other traits constant) and simulate their labor market outcomes. As seen in the table, a one standard deviation increase in years of education (approx. 2.8 years) increases wages by 16-20 percent for both men and women. It also reduces unemployment and job spell lengths, particularly for women. A shorter job spell is not necessarily detrimental if the individual changes jobs because of good outside offers. An increase in the cognitive

Table 5: Estimated index coefficients associated with characteristics (education, cognitive ability, personality traits, cohort) by gender†

	$\log a$		$\log \lambda_U$		$\log \eta$		$\log\left(\frac{\alpha}{1-\alpha}\right)$	
	Male	Female	Male	Female	Male	Female	Male	Female
Constant	3.034	2.920	-1.751	-1.751	-6.164	-6.051	-0.074	-0.295
	(0.011)	(0.010)	(0.017)	(0.012)	(0.047)	(0.039)	(0.024)	(0.023)
Education	0.081	0.080	0.244	0.244	-0.684	-0.398	0.146	0.131
	(0.006)	(0.010)	(0.011)	(0.012)	(0.036)	(0.052)	(0.015)	(0.023)
Cognitive ability	0.019	0.030	0.070	0.189	0.141	0.092	0.005	-0.042
	(0.007)	(0.005)	(0.006)	(0.008)	(0.041)	(0.024)	(0.018)	(0.007)
Openness	-0.001	0.013	0.039	0.029	0.099	-0.076	-0.043	-0.069
	(0.007)	(0.008)	(0.011)	(0.010)	(0.044)	(0.016)	(0.019)	(0.015)
Conscientiousness	0.043	0.012	0.080	0.089	-0.287	-0.294	0.007	0.065
	(0.008)	(0.010)	(0.012)	(0.005)	(0.039)	(0.054)	(0.022)	(0.027)
Extraversion	0.004	0.008	0.033	0.010	0.092	0.026	-0.038	-0.024
	(0.007)	(0.009)	(0.012)	(0.011)	(0.050)	(0.050)	(0.022)	(0.021)
Agreeableness	-0.026	-0.004	-0.024	-0.090	0.116	0.115	-0.005	-0.150
	(0.007)	(0.010)	(0.010)	(0.009)	(0.021)	(0.052)	(0.019)	(0.023)
Emotional stability	0.025	0.010	0.024	0.114	-0.143	-0.108	0.013	0.132
	(0.008)	(0.012)	(0.011)	(0.011)	(0.047)	(0.038)	(0.023)	(0.030)
Cohort (ref group: 37-48)								
25-36	-0.124	-0.034	0.130	-0.050	0.582	0.541	0.018	-0.043
	(0.013)	(0.014)	(0.023)	(0.019)	(0.065)	(0.093)	(0.032)	(0.026)
49-60	-0.048	-0.044	-0.517	-0.378	-0.308	-0.262	-0.004	-0.072
	(0.023)	(0.014)	(0.025)	(0.025)	(0.101)	(0.110)	(0.064)	(0.030)

†This table reports estimated parameter coefficients for the heterogeneous specification. Asymptotic standard errors are reported in parentheses.

test score has similar effects to an increase in years of education, although the magnitudes of the wage effects are smaller and the effects on job spell lengths larger.

An increase in conscientiousness affects men and women in similar ways, increasing their wages, reducing unemployment spell lengths and increasing job spell lengths. Similarly, increasing emotional stability increases wages and decreases unemployment spell lengths. It also increases job spell length for men but decreases it for women. Increasing openness to experience has almost no effect on wages but decreases the unemployment spell lengths for both men and women and reduces the employment spell length for men. A ceteris paribus increase in extraversion does not impact wages much, but it decreases unemployment and job spell lengths, with larger effects for men. Lastly, agreeableness mostly negatively impacts labor market outcomes. For men, an increase in agreeableness leads to lower wages and increased unemployment spell lengths. For women, agreeableness also lowers wages, increases unemployment spells and increases job spell length. In summary, we find both cognitive and noncognitive traits to be important determinants of wages and labor market

Table 6: Effects of 1SD changes in cognitive and non-cognitive traits on labor market outcomes†

	Average wage		Unemp. spell		Job spell	
	Men	Women	Men	Women	Men	Women
Baseline	15.33	12.52	7.30	7.43	71.96	55.68
Education (+1 SD)	19.8%	16.3%	-21.5%	-21.6%	-3.6%	-10.4%
Cognitive ability (+1 SD)	2.6%	2.5%	-6.7%	-17.2%	-8.0%	-15.1%
Openness (+1 SD)	-1.4%	-1.2%	-3.8%	-2.8%	-5.0%	-0.6%
Conscientiousness (+1 SD)	6.8%	5.3%	-7.6%	-8.5%	1.3%	0.0%
Extraversion (+1 SD)	-0.8%	-0.1%	-3.2%	-1.0%	-4.5%	-1.4%
Agreeableness (+1 SD)	-3.1%	-7.1%	2.4%	9.4%	-0.8%	3.9%
Emotional stability (+1 SD)	4.0%	7.4%	-2.3%	-10.8%	2.3%	-5.8%

†The first row shows labor market outcome values in steady-state under the baseline model. Rows (2)-(6) show the deviation from baseline outcomes implied by a ceteris paribus one standard deviation increase in each trait.

dynamics. Of the Big Five personality traits, conscientiousness and emotional stability are the key traits contributing to higher wages and more stable employment. Agreeableness substantially lowers wages for both men and women.

There are various reasons why personality traits could be important determinants of labor market outcomes. As seen in Table 5, some traits directly enhance worker's productive ability. People who are more conscientiousness tend to be well organized, dependable and hard-working, which are all characteristics associated with more productive workers (Barrick and Mount (1991); Salgado (1997); Hurtz and Donovan (2000); Cubel et al. (2016)). Other traits operate through different channels. For example, individuals with higher emotional stability and lower agreeableness may be more likely to negotiate pay rises. Evdokimov and Rahman (2014) provide experimental evidence that managers allocate less money to more agreeable workers. Although previous papers also find associations between personality traits and wages (Mueller and Plug (2006); Heineck and Anger (2010); Risse et al. (2018)), the mechanisms have not been explored.[41]

Table 7 shows the contribution of each observed trait to wages in percentages by model channels. Education increases wages through all model channels, with ability being quantitatively the most important one. The cognitive score also affects wages primarily through its effect on ability. The five personality traits operate through multiple model channels, with ability and surplus division having the greatest quantitative significance. Two traits -

[41] Our estimates are mostly consistent with the literature exploring the gender-specific association between wages and personality traits. For example, Nyhus and Pons (2005) note that emotional stability is positively associated with wages for both women and men, while agreeableness is associated with lower wages for women. Using GSOEP data, Braakmann (2009) finds agreeableness, conscientiousness and neurotism matter for both wages and employment.

Table 7: Decomposing the effects of observed traits on wages by model channel†

		All channels	Ability	Surplus division	Transition prob
Education (+1 SD)	M	19.8%	8.4%	5.0%	5.4%
	F	16.3%	8.3%	5.1%	2.3%
Cognitive ability (+1 SD)	M	2.6%	1.9%	0.2%	0.5%
	F	2.5%	3.0%	-1.7%	1.2%
Openness (+1 SD)	M	-1.4%	-0.1%	-1.5%	0.2%
	F	-1.2%	1.3%	-2.7%	0.3%
Conscientiousness (+1 SD)	M	6.8%	4.4%	0.2%	2.0%
	F	5.3%	1.2%	2.5%	1.5%
Extraversion (+1 SD)	M	-0.8%	0.4%	-1.3%	0.2%
	F	-0.1%	0.8%	-0.9%	0.1%
Agreeableness (+1 SD)	M	-3.1%	-2.6%	-0.2%	-0.4%
	F	-7.1%	-0.4%	-5.9%	-0.7%
Emotional stability (+1 SD)	M	4.0%	2.5%	0.4%	1.0%
	F	7.4%	1.0%	5.1%	1.1%

†The table shows the ceteris paribus effect of a one standard deviation (SD) increase in each of the traits.

emotional stability and conscientiousness- have the largest positive effect on wages. In contrast, Agreeableness negatively impacts wages. Although the overall effects are qualitatively consistent between men and women, the primary model channel differs. For men, ability (productivity) is the primary channel and for women bargaining/surplus division is the primary channel. The net wage effects of extraversion and openness to experience are small; the positive effects of extraversion on ability (productivity) are offset by negative effects on surplus division.

6.2 Understanding gender wage gap using an extended Oaxaca-Blinder decomposition

The Oaxaca (1973) and Blinder (1973) decomposition approach is often used in linear model settings to analyze the determinants of gender or racial wage gaps. In this section, we adapt the decomposition method to our nonlinear setting to analyze which model channels contribute the most to gender wage gaps. To generate the results in Table 8, we simulate outcomes under the heterogeneous specification in three different ways. First, we simulate outcomes under the baseline model with the parameter vectors that were previously estimated. Let γ_m and γ_f denote the male and female estimated parameter values and let z_m and z_f denote the male and female characteristics (as in the data). $w(\gamma_m, z_m)$ and $w(\gamma_f, z_f)$ denote the average simulated wages for males and females. The average wage gap is $w(\gamma_m, z_m) - w(\gamma_f, z_f)$. Second, we perform a simulation where we adjust the mean female traits (upward or downward with additive constants) to be equal to the mean for males, but keeping their parameter values as estimated.[42] We denote the average wages thus obtained by $w(\gamma_f, z_m)$. Third, we do a simulation were we keep female traits as observed in the data but give females the male estimated parameter values, denoting the average wages thus obtained by $w(\gamma_m, z_f)$. The percentage of the average wage gap attributable to females having different trait levels from males is

$$\Delta^{trait} = \frac{w(\gamma_f, z_f) - w(\gamma_f, z_m)}{w(\gamma_m, z_m) - w(\gamma_f, z_f)}$$

The percentage of the wage gap attributable to females having different parameter values from males is:

$$\Delta^{coef} = \frac{w(\gamma_f, z_f) - w(\gamma_m, z_f)}{w(\gamma_m, z_m) - w(\gamma_f, z_f)}$$

We perform this decomposition overall and for subsets of parameter values that enter through different model channels. We also perform the decomposition separately for the different cognitive and noncognitive traits.

As seen in Table 8, education and cognitive ability do not account for the wage gap. If we simulate labor market outcomes with the adjusted female traits but with the original female parameters values, we find that the average wage gap would be 4.3 percent larger than it currently is (that is, the wage gap would actually increase). Similarly, simulating outcomes with female characteristics but using the parameters estimated for males also increases the wage gap, albeit to a lesser extent (0.2 percent). Cognitive ability differences, either in levels or in terms of associated estimated parameter values, have little effect on the wage gap.

Table 8 (row 5) shows that the gender wage gap is largely explained by differences in male-female personality trait levels. If we simulate labor market outcomes with the adjusted female traits and female estimated parameters, the wage gap is reduced by 19.6 percent. Looking at the last three columns, it is clear that the surplus division model channel accounts for the most of the wage gap. That is, females have personality traits that lead them to have lower bargaining power.

Examining the personality traits separately, we see that two traits largely explain the gender wage gap: agreeableness and emotional stability. As was seen in Table 1, these traits

[42] The additive constant to equate the means is $\bar{z}_m - \bar{z}_f$.

differ substantially, on average, for men and women. As seen in Table 6, agreeableness is negatively remunerated while emotional stability is positively remunerated. The fact that women have on average higher levels of agreeableness and lower levels of emotional stability leads to a significant wage disadvantage. The gender gap explained by endowment differences in agreeableness and emotional stability is 9.5 percent and 13.6 percent. For these two traits, parameter value differences contribute to the gender wage gap, although the magnitudes are much smaller. Partly offsetting these effects is the finding that women display on average greater conscientiousness than men—a trait that is positive remunerated. The different conscientiousness values narrow the gender wage gap by 3.6 percent. In general, gender differences in personality endowments have a stronger quantifiable role in explaining the gender pay gap than do gender differences in the return to personality traits.

Thus, we find that gender wage disparities are primarily attributable to gender differences in personality traits. Two traits in particular, emotional stability and agreeableness, put women at a disadvantage in the labor market. An important mechanism through which these traits affect labor market outcomes is by reducing women's bargaining power relative to men.

Table 8: Decomposition of the gender wage-gap†

	Counterfactual	All channels	Ability	Surplus division	Transition
Education	Δ^{trait}	4.3%	2.6%	1.7%	-0.0%
	Δ^{coef}	0.2%	-0.1%	-0.3%	-0.5%
Cognitive ability	Δ^{trait}	-0.4%	-0.4%	0.3%	-0.2%
	Δ^{coef}	-0.7%	0.2%	-0.3%	-0.8%
Big Five personality traits	Δ^{trait}	-19.6%	0.6%	-18.7%	-1.4%
	Δ^{coef}	-6.4%	0.5%	-6.9%	-1.0%
Openness to experience	Δ^{trait}	-1.2%	1.1%	-2.3%	-0.0%
	Δ^{coef}	0.1%	0.6%	-0.6%	-0.0%
Conscientiousness	Δ^{trait}	3.6%	1.1%	2.4%	0.3%
	Δ^{coef}	-0.5%	-2.2%	1.3%	-0.0%
Extraversion	Δ^{trait}	-0.4%	0.9%	-1.1%	-0.2%
	Δ^{coef}	-0.2%	0.3%	0.4%	-0.9%
Agreeableness	Δ^{trait}	-9.5%	-0.6%	-8.1%	-0.7%
	Δ^{coef}	-2.5%	1.0%	-4.0%	-0.0%
Emotional stability	Δ^{trait}	-13.6%	-1.9%	-9.6%	-2.0%
	Δ^{coef}	-3.2%	0.7%	-3.8%	-0.4%

† Rows Δ^{trait} capture the change in gender wage gap when increasing females' traits by a constant value $(\bar{z}_m - \bar{z}_f)$ that equates the means to those of males and keeping the parameter values the same as estimated. The numbers in these rows are calculated as $\frac{w(\gamma_f,z_f)-w(\gamma_f,z_m)}{w(\gamma_m,z_m)-w(\gamma_f,z_f)}$. Rows Δ^{coef} capture the changes in gender wage gap when giving female workers the male parameter values but keeping their trait levels unchanged. The numbers in these rows are calculated as $\frac{w(\gamma_f,z_f)-w(\gamma_m,z_f)}{w(\gamma_m,z_m)-w(\gamma_f,z_f)}$.

6.3 Exploring effects of interventions aimed at modifying personality traits

There exists a clinical psychology literature that examines whether and to what extent personality traits are malleable in response to clinical interventions, such as cognitive-behavioral therapy or pharmacological treatments(e.g., Barlow et al. (2014); Quilty et al. (2014); Soskin et al. (2012)). Such interventions are often targeted at individuals diagnosed with mental health problems, such as avoidant personality disorder, social anxiety disorder, depression, or eating disorders (which are comorbid with depression and anxiety). This literature finds that even relatively short-term interventions (6-8 weeks) can lead to lasting personality trait changes.

A recent meta analysis by Roberts et al. (2017) summarizes results of 207 studies of the effects of clinical interventions. Table 9 shows the range of the effect sizes (ES) that Roberts et al. (2017) report, based on the full set of studies considered and on the subset of experimental (RCT) studies. As seen in the table, the interventions affect multiple personality traits, but have the greatest impact on emotional stability. Roberts et al. (2017) also examine how estimated effect sizes depend on treatment intervention duration. They conclude that a minimum of 4 weeks is needed to see significant intervention effects but that there is little marginal benefit from going beyond 8 weeks. They argue that the optimal intervention duration is in the 6-8 week range.

Table 9: Personality trait effect size estimates†

Moderator	Full Sample ES [95% CI]	Experimental Studies ES [95% CI]
Extraversion	.23 [.17, .29]	.38 [.18, .58]
Agreeableness	.15 [.11, .20]	.23 [.08, .38]
Conscientiousness	.19 [.14, .23]	.06 [-.05, .16]
Emotional stability	.57 [.52, .62]	.69 [.45, .93]
Openness	.13 [.07, .18]	.36 [.23, .49]

†Reported numbers are based on Tables 2 and 3 in Roberts et al. (2017). The third column reports results based on experimental studies only. ES=effect size; CI=confidence interval.

Using the effect size from the full sample estimates shown in the Table 9, we examine the potential for such interventions to impact workers' labor market outcomes. In our simulations, we target the intervention at individuals with relatively low emotional stability scores, those in the bottom 10th percentile, 25th percentile, and 50th percentile of the distribution.[43] Women tend to be overrepresented in the lower tail, which is also consistent with

[43] According to the German Association for Psychiatry, Psychotherapy, and Psychosomatics, around 1 in 3

Table 10: The simulated effects of intervention on gender personality trait gaps

	Baseline Sample	Target group: individuals below emotional stability quantile X		
		10 quantile	25 quantile	50 quantile
Percent female	-	71.4%	65.3%	59.6%
Openness to experience	-0.195	-0.205	-0.213	-0.218
Conscientiousness	-0.215	-0.231	-0.243	-0.249
Extraversion	-0.274	-0.293	-0.307	-0.315
Agreeableness	-0.320	-0.332	-0.342	-0.346
Emotional stability	0.435	0.390	0.354	0.335

their overpresentation among those treated clinically. For example, of the 20,024 individuals in the 207 studies considered by Roberts et al. (2017), female participants accounted for 63.41 percent of the cases.

The effects of a hypothetical mental health intervention targeted at different size groups on average gender gaps in personality traits are reported in Table 10. The gap in emotional stability shrinks whereas the gaps in the other four personality traits widen. As increasing emotional stability raises wages and increasing agreeableness lowers wages, the net effect of the intervention on the gender wage gap needs to be quantitatively assessed. Table 11 shows how intervening with different size target groups affects wage inequality and the gender wage gap. When 10 percent of individuals receive the intervention, inequality, as measured by the 90-10 wage ratio, reduces by 0.1 percent and the median gender wage gap reduces by 2.6 percent. When a larger population fraction is targeted (25 percent), there is a 1.3 percent reduction in the 90-10 wage ratio and a 3.1 percent reduction in the gender wage gap for the median worker. The largest gender wage gap reductions occur at the lower quantiles of the wage distribution. These results suggest that mental health interventions can be an effective means of increasing the wages of individuals whose personality traits put them at a labor market disadvantage, which, in turn, also decreases wage inequality and the gender wage gap. Increasing accessibility to mental health services is potentially an effective policy for improving labor market outcomes.

German adults suffer from some mental health issue every year, with anxiety disorders and mood disorders being the most common. However, only 18.9 percent of these people seek assistance from health service providers.

Table 11: Wage gaps with intervention targeted at different size groups†

	Wage in the population				Gender wage gap at n-th percentile		
	Mean	90-10 ratio	Male	Female	25 percent	50 percent	75 percent
Baseline	13.90	3.55	15.33	12.52	0.187	0.179	0.182
Counterfactual experiments							
10%	13.94	3.54	15.35	12.58	0.182	0.174	0.181
	0.3%	-0.1%	0.1%	0.5%	-3.7%	-2.6%	-0.9%
25%	14.02	3.50	15.40	12.67	0.178	0.173	0.175
	0.8%	-1.3%	0.5%	1.2%	-4.9%	-3.1%	-3.9%
50%	14.14	3.47	15.52	12.80	0.173	0.170	0.175
	1.7%	-2.2%	1.2%	2.3%	-7.8%	-4.9%	-4.0%

†The gender wage gap at a given $n-th$ percentile is calculated by $\frac{w_m^n - w_f^n}{w_m^n}$, where w_m^n is the wage rate at $n-th$ percentile of male's wage distribution while w_f^n is the wage rate at $n-th$ percentile of female's wage distribution.

7 Conclusion

This paper extends a canonical job search model to incorporate a rich set of individual characteristics, including both cognitive and noncognitive attributes. We use the estimated model to explore the determinants of gender wage gaps and to analyze the effects of potential policy interventions on gender wage disparities and overall inequality. Our analysis considered two nested model specifications that differ in the degree of parameter heterogeneity. Likelihood ratio tests and goodness of fit criteria support the use of the model allowing for greater heterogeneity. The model estimates show that education, cognitive ability and personality traits are important determinants of productivity, bargaining and job offer arrival rates for both men and women.

We use the estimated model to simulate steady state labor market outcomes. Two personality traits, conscientiousness and emotional stability, contribute to more favorable labor market outcomes for both men and women. Higher values of these traits lead to higher wages and more stable employment. One trait, agreeableness, systematically leads to worse labor market outcomes.

We developed a Oaxaca-Blinder type decomposition, extended to our nonlinear model setting, to analyze the contribution of different individual traits and model channels in accounting for the gender wage gap. The results showed that education and cognitive ability do not contribute to gender wage disparities. In fact, gender differences in education levels and in the returns to education (in terms of productivity and bargaining) tend to reduce the gender wage gap. Similarly, gender differences in cognitive ability do not explain wage gaps.

Personality traits emerge as the primary factors accounting for gender wage gaps, particularly as they operate through the surplus division (bargaining) channel of the job search model. Our analysis reveals that women have substantially lower bargaining power than men, largely because they have higher average levels of agreeableness and lower levels of emotional stability. These two traits also serve to reduce wages through the ability and job transition model channels. There are some gender differences in the trait valuations that exacerbate wage gaps, but the gender differences in trait levels account for the vast majority of the wage gap. The wage gap would be reduced by 19.6 percent (from 18.9 to 15.2) if women had the same average personality trait levels as men. Our evidence adds to the growing body of literature demonstrating that noncognitive attributes, such as personality traits, are important determinants of adult labor market success and gender inequality.

Lastly, we used the estimated job search model to study the potential effects of mental health interventions that have been shown in the clinical psychology literature to modify some aspects of personality. Using effect size estimates in the range found in previous experimental studies, we show that an intervention targeted at individuals with low emotional stability scores has the potential to improve their wage outcomes and shrink the average gender wage gap by 2-6 percent. The findings suggest that increasing access and take-up of mental health services could be a way of enhancing workers' labor market outcomes and reducing inequalities.

References

Almlund, M., Duckworth, A. L., Heckman, J., and Kautz, T. (2011). Personality psychology and economics. In *Handbook of the Economics of Education*, volume 4, pages 1–181. Elsevier.

Amano-Patino, N., Baron, T., and Xiao, P. (2020). Equilibrium wage-setting and the life-cycle gender pay gap. *Cambridge Working Papers in Economics 2010, Faculty of Economics, University of Cambridge.*

Babcock, L., Laschever, S., Gelfand, M., and Small, D. (2003). Nice girls don't ask. *Harvard Business Review*, 81(10):14–16.

Bagby, R. M., Quilty, L. C., Segal, Z. V., McBride, C. C., Kennedy, S. H., and Costa Jr,

P. T. (2008). Personality and differential treatment response in major depression: a randomized controlled trial comparing cognitive-behavioural therapy and pharmacotherapy. *The Canadian Journal of Psychiatry*, 53(6):361–370.

Barlow, D. H., Sauer-Zavala, S., Carl, J. R., Bullis, J. R., and Ellard, K. K. (2014). The nature, diagnosis, and treatment of neuroticism: Back to the future. *Clinical Psychological Science*, 2(3):344–365.

Barrick, M. R. and Mount, M. K. (1991). The big five personality dimensions and job performance: a meta-analysis. *Personnel psychology*, 44(1):1–26.

Bartolucci, C. (2013). Gender wage gaps reconsidered a structural approach using matched employer-employee data. *Journal of Human Resources*, 48(4):998–1034.

Blau, F. D. and Kahn, L. M. (2000). Gender differences in pay. *Journal of Economic Perspectives*, 14(4):75–99.

Blau, F. D. and Kahn, L. M. (2006). The US gender pay gap in the 1990s: Slowing convergence. *ILR Review*, 60(1):45–66.

Bleidorn, W., Hopwood, C. J., and Lucas, R. E. (2018). Life events and personality trait change. *Journal of personality*, 86(1):83–96.

Blinder, A. S. (1973). Wage discrimination: reduced form and structural estimates. *Journal of Human resources*, pages 436–455.

Borghans, L., Duckworth, A. L., Heckman, J. J., and Ter Weel, B. (2008). The economics and psychology of personality traits. *Journal of Human Resources*, 43(4):972–1059.

Bound, J., Brown, C., Duncan, G. J., and Rodgers, W. L. (1994). Evidence on the validity of cross-sectional and longitudinal labor market data. *Journal of Labor Economics*, 12(3):345–368.

Bound, J. and Krueger, A. B. (1991). The extent of measurement error in longitudinal earnings data: Do two wrongs make a right? *Journal of labor economics*, 9(1):1–24.

Bowlus, A. J. (1997). A search interpretation of male-female wage differentials. *Journal of Labor Economics*, 15(4):625–657.

Bowlus, A. J. and Grogan, L. (2008). Gender wage differentials, job search, and part-time employment in the uk. *Oxford Economic Papers*, 61(2):275–303.

Braakmann, N. (2009). The role of psychological traits for the gender gap in full-time employment and wages: evidence from germany. *SOEP Paper*, (162).

Cahuc, P., Postel-Vinay, F., and Robin, J.-M. (2006). Wage bargaining with on-the-job search: Theory and evidence. *Econometrica*, 74(2):323–364.

Caliendo, M., Cobb-Clark, D. A., and Uhlendorff, A. (2015). Locus of control and job search strategies. *Review of Economics and Statistics*, 97(1):88–103.

Card, D., Cardoso, A. R., and Kline, P. (2015). Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women. *The Quarterly Journal of Economics*, 131(2):633–686.

Cattan, S. (2013). Psychological traits and the gender wage gap. *The Institute for Fiscal Studies*.

Cobb-Clark, D. A. and Schurer, S. (2012). The stability of big-five personality traits. *Economics Letters*, 115(1):11–15.

Cobb-Clark, D. A. and Schurer, S. (2013). Two economists' musings on the stability of locus of control. *The Economic Journal*, 123(570):F358–F400.

Collischon, M. (2021). Personality traits as a partial explanation for gender wage gaps and glass ceilings. *Research in Social Stratification and Mobility*, 73:100596.

Costa Jr, P. T. and McCrae, R. R. (1988). From catalog to classification: Murray's needs and the five-factor model. *Journal of personality and social psychology*, 55(2):258.

Croson, R. and Gneezy, U. (2009). Gender differences in preferences. *Journal of Economic Literature*, 47(2):448–74.

Cubel, M., Nuevo-Chiquero, A., Sanchez-Pages, S., and Vidal-Fernandez, M. (2016). Do personality traits affect productivity? evidence from the laboratory. *The Economic Journal*, 126(592):654–681.

Dey, M. and Flinn, C. (2008). Household search and health insurance coverage. *Journal of Econometrics*, 145(1-2):43–63.

Dey, M. S. and Flinn, C. J. (2005). An equilibrium model of health insurance provision and wage determination. *Econometrica*, 73(2):571–627.

Dittrich, M., Knabe, A., and Leipold, K. (2014). Gender differences in experimental wage negotiations. *Economic Inquiry*, 52(2):862–873.

Dohmen, T., Falk, A., Huffman, D., Sunde, U., Schupp, J., and Wagner, G. G. (2011). Individual risk attitudes: Measurement, determinants, and behavioral consequences. *Journal of the European Economic Association*, 9(3):522–550.

Evdokimov, P. and Rahman, D. (2014). Personality and bargaining power. Technical report, Mimeo, University of Minnesota.

Flabbi, L. (2010a). Gender discrimination estimation in a search model with matching and bargaining. *International Economic Review*, 51(3):745–783.

Flabbi, L. (2010b). Prejudice and gender differentials in the US labor market in the last twenty years. *Journal of Econometrics*, 156(1):190–200.

Flinn, C., Gemici, A., and Laufer, S. (2017). Search, matching and training. *Review of Economic Dynamics*, 25:260–297.

Flinn, C. and Heckman, J. (1982). New methods for analyzing structural models of labor force dynamics. *Journal of Econometrics*, 18(1):115–168.

Flinn, C. and Mullins, J. (2015). Labor market search and schooling investment. *International Economic Review*, 56(2):359–398.

Flinn, C. J. (2002). Labour market structure and inequality: A comparison of italy and the us. *The Review of Economic Studies*, 69(3):611–645.

Flinn, C. J. (2006). Minimum wage effects on labor market outcomes under search, matching, and endogenous contact rates. *Econometrica*, 74(4):1013–1062.

Flinn, C. J. (2011). *The minimum wage and labor market outcomes*. MIT press.

Flinn, C. J., Todd, P. E., and Zhang, W. (2018). Personality traits, intra-household allocation and the gender wage gap. *European Economic Review*, 109:191–220.

Fraley, R. C. and Roberts, B. W. (2005). Patterns of continuity: a dynamic model for conceptualizing the stability of individual differences in psychological constructs across the life course. *Psychological review*, 112(1):60.

Gottfries, A. (2018). Bargaining with renegotiation in models with on-the-job search. *Available at SSRN: https://ssrn.com/abstract=3058870 or http://dx.doi.org/10.2139/ssrn.3058870*.

Heckman, J. J. and Raut, L. K. (2016). Intergenerational long-term effects of preschool-structural estimates from a discrete dynamic programming model. *Journal of Econometrics*, 191(1):164–175.

Heineck, G. (2011). Does it pay to be nice? personality and earnings in the United Kingdom. *ILR Review*, 64(5):1020–1038.

Heineck, G. and Anger, S. (2010). The returns to cognitive abilities and personality traits in germany. *Labour economics*, 17(3):535–546.

Hurtz, G. M. and Donovan, J. J. (2000). Personality and job performance: The big five revisited. *Journal of applied psychology*, 85(6):869.

Lavy, V. (2013). Gender differences in market competitiveness in a real workplace: Evidence from performance-based pay tournaments among teachers. *The Economic Journal*, 123(569):540–573.

Lindenberger, U. and Baltes, P. B. (1995). Testing-the-limits and experimental simulation: Two methods to explicate the role of learning in development. *Human Development*, 38(6):349–360.

Liu, K. (2016). Explaining the gender wage gap: Estimates from a dynamic model of job changes and hours changes. *Quantitative Economics*, 7(2):411–447.

Lüdtke, O., Roberts, B. W., Trautwein, U., and Nagy, G. (2011). A random walk down university avenue: life paths, life events, and personality trait change at the transition to university life. *Journal of personality and social psychology*, 101(3):620.

Manning, A. and Saidi, F. (2010). Understanding the gender pay gap: what's competition got to do with it? *ILR Review*, 63(4):681–698.

McCrae, R. R. and Costa Jr, P. T. (1994). The stability of personality: Observations and evaluations. *Current directions in psychological science*, 3(6):173–175.

McCrae, R. R., Costa Jr, P. T., Ostendorf, F., Angleitner, A., Hřebíčková, M., Avia, M. D., Sanz, J., Sanchez-Bernardos, M. L., Kusdil, M. E., Woodfield, R., et al. (2000). Nature over nurture: temperament, personality, and life span development. *Journal of personality and social psychology*, 78(1):173.

McGee, A. D. (2015). How the perception of control influences unemployed job search. *ILR Review*, 68(1):184–211.

Morchio, I. and Moser, C. (2020). The gender pay gap: Micro sources and macro consequences. *Available at SSRN: https://ssrn.com/abstract=3176868 or http://dx.doi.org/10.2139/ssrn.3176868*.

Mortensen, D. T. (2005). Wage dispersion: Why are similar workers paid differently? *MIT Press Books*.

Moscarini, G. (2008). Job-to-job quits and corporate culture: New results. In *2008 Meeting Papers*, volume 153. Society for Economic Dynamics.

Mueller, G. and Plug, E. (2006). Estimating the effect of personality on male and female earnings. *ILR Review*, 60(1):3–22.

Nyhus, E. K. and Pons, E. (2005). The effects of personality on earnings. *Journal of Economic Psychology*, 26(3):363–384.

Nyhus, E. K. and Pons, E. (2012). Personality and the gender wage gap. *Applied Economics*, 44(1):105–118.

Oaxaca, R. (1973). Male-female wage differentials in urban labor markets. *International economic review*, pages 693–709.

Pitman, E. J. (1939). The estimation of the location and scale parameters of a continuous population of any given form. *Biometrika*, 30(3/4):391–421.

Postel-Vinay, F. and Robin, J.-M. (2002). Equilibrium wage dispersion with worker and employer heterogeneity. *Econometrica*, 70(6):2295–2350.

Quilty, L. C., Dozois, D. J., Lobo, D. S., Ravindran, L. N., and Bagby, R. M. (2014). Cognitive structure and processing during cognitive behavioral therapy vs. pharmacotherapy for depression. *International Journal of Cognitive Therapy*, 7(3):235–250.

Risse, L., Farrell, L., and Fry, T. R. (2018). Personality and pay: do gender gaps in confidence explain gender gaps in wages? *Oxford Economic Papers*, 70(4):919–949.

Roberts, B. W., Luo, J., Briley, D. A., Chow, P. I., Su, R., and Hill, P. L. (2017). A systematic review of personality trait change through intervention. *Psychological Bulletin*, 143(2):117.

Roberts, B. W., Wood, D., and Caspi, A. (2008). The development of personality traits in adulthood.

Romeo, C. J. (2001). Controlling for seam problems in duration model estimates: With application to the current population survey and the computer aided telephone interview/computer aided personal interview overlap survey. *Journal of Human Resources*, pages 467–499.

Salgado, J. F. (1997). The five factor model of personality and job performance in the european community. *Journal of Applied psychology*, 82(1):30.

Säve-Söderbergh, J. (2007). Are women asking for low wages? gender differences in wage bargaining strategies and ensuing bargaining success. *Working paper: Institutet for Social Forskning Volumes 7-2007 of Working papers / Swedish Institute for Social Research.*

Semykina, A. and Linz, S. J. (2007). Gender differences in personality and earnings: Evidence from russia. *Journal of Economic Psychology*, 28(3):387–410.

Shimer, R. (2006). On-the-job search and strategic bargaining. *European Economic Review*, 50(4):811–830.

Soskin, D. P., Carl, J. R., Alpert, J., and Fava, M. (2012). Antidepressant effects on emotional temperament: toward a biobehavioral research paradigm for major depressive disorder. *CNS Neuroscience & Therapeutics*, 18(6):441–451.

Stuhlmacher, A. F. and Walters, A. E. (1999). Gender differences in negotiation outcome: A meta-analysis. *Personnel Psychology*, 52(3):653–677.

Todd, P. E. and Zhang, W. (2020). A dynamic model of personality, schooling, and occupational choice. *Quantitative Economics*, 11(1):231–275.

Wolpin, K. I. (1987). Estimating a structural search model: The transition from school to work. *Econometrica*, 55(4):801–17.

Xiao, P. (2020). Wage and employment discrimination by gender in labor market equilibrium. *Job market paper, Yale university.*

A Appendices

A.1 Model Solutions

A.1.1 Solving the reservation match value θ^* and value of unemployment V_U when $a = 1$

In this section, we describe how to numerically solve V_U and θ^* for individuals with "standard" ability $a = 1$. The basic algorithm is described in Chapter 10, Flinn (2011). It solves a fixed point problem in V_U in a recursive manner: (1) we begin with a initial guess of the value of employment $\tilde{V}_U$; (2) we then solve the optimal reservation match value $\hat{\theta}^*$ and its associated value of unemployment $\hat{V}_U$; (3) we replace $\tilde{V}_U$ with $\hat{V}_U$ and repeat step 1 and 2 until $\tilde{V}_U$ and $\hat{V}_U$ converges. Below, we describe the solution method step by step.

We first discretize the continuous θ interval into L grid points $\{\theta_1, ..., \theta_L\}$ with probability $\{p_1, ..., p_L\}$. To initialize the algorithm, we set a initial guess of unemployment $\tilde{V}_U$ to be equal to the flow utility ab. The recursive algorithm solves θ^* and V_U by the following steps:

1. Solve for the value of employment $V_E(\theta_L)$ and bargained wages $w(\theta_L)$ at the largest match value θ_L. A job with matching value θ_L would have no further job-to-job mobility. The only way such a job can end is through exogenous termination, which occurs at the rate η.

$$V_E(\theta_L) = \frac{w(\theta_L) + \eta \tilde{V}_U}{\rho + \eta}$$

with the wage

$$w(\theta_L) = a\left(\alpha\theta_L + (1-\alpha)\rho\tilde{V}_U\right)$$

Given the value of $V_E(\theta_L)$ and $w(\theta_L)$, we can calculate the (potential) value of unemployment when assuming the reservation match value $\theta^* = \theta^L$

$$\bar{V}_U(\theta^* = \theta_L) = \frac{ab + \lambda_U p_L V_E(\theta_L)}{\rho + \lambda_U p_L}$$

2. Sequentially solve for $V_E(\theta_l)$ and $w(\theta_l)$ as well as $\bar{V}_U(\theta^* = \theta_l)$ at match value grids θ_l in a backward order. Given $(V_E(\theta_{l+1}), ..., V_E(\theta_L))$, we can solve for the wage $w(\theta_l)$ associated with the match value $w(\theta_l)$ as

$$w(\theta_l) = a\left(\alpha\theta_l + (1-\alpha)\left(\left(\rho + \lambda_E p_{l+1}^+\right)\tilde{V}_U - \lambda_E \sum_{i\geq l+1}^{L} p_i V_E(\theta_i))\right)\right).$$

The value of employment associated with match value θ_l is given by

$$V_E(\theta_l) = \frac{w(\theta_l) + \eta\tilde{V}_U + \lambda_E \sum_{i\geq l}^{L} p_i V_E(\theta_i)}{\rho + \eta + \lambda_E p_l^+}$$

where the notation $p_l^+ = \sum_{i\geq l}^{L} p_i$.

Given the value of $V_E(\theta_l)$ and $w(\theta_l)$, we can calculate the (potential) value of unemployment when assuming the reservation match value equals to $\theta^* = \theta^l$

$$\bar{V}_U(\theta^* = \theta_l) = \frac{ab + \lambda_U \sum_{i\geq l}^{L} p_i V_E(\theta_i)}{\rho + \lambda_U p_l^+}$$

3. Determine the "optimal" reservation match quality θ^*. For all match quality $\{\theta_1, ..., \theta_L\}$, each "potential" reservation match value implies a value of unemployment given by $\bar{V}_U(\theta^* = \theta_l)$. The "optimal" reservation match value is the one that produces that highest value of unemployment , i.e.,

$$j = \arg\max_l \left\{\bar{V}_U(\theta^* = \theta_l)\right\}_{l=1}^{L}$$
$$\bar{V}_U^{new} = \bar{V}_U(\theta^* = \theta_j), \theta^* = \theta_j$$

4. Stop if $\bar{V}_U^{new} = \tilde{V}_U$. Otherwise update $\tilde{V}_U$ with the new value $\bar{V}_U^{new}$ and repeat step 1 to 3 until $\bar{V}_U^{new}$ and $\tilde{V}_U$ are convergent.

A.1.2 The likelihood function

Our model is estimated using maximum likelihood. We describe in detail how the likelihood function of an employment cycle is constructed in this section. As previously noted, we classify the employment cycles into two categories based on worker's employment status at the beginning of the employment cycle. If the employment cycle starts with an unemployment spell, then the relevant variables in the employment cycle are

$$\text{Employment cycle} = \underbrace{\{t_U, r_U\}}_{\text{Unemployment spell}}, \underbrace{\{\tilde{w}_m, t_m, r_m, q_m\}_{m=1}^{M}}_{\text{Consecutive M jobs}}$$

On the other hand, the relevant variables included in the employment cycle are

$$\text{Employment cycle} = \underbrace{\{\tilde{w}_m, t_m, r_m, q_m\}_{m=1}^{M}}_{\text{Consecutive M jobs}}$$

For the unemployment spell, t_U is the length of the unemployment spell, r_U is whether the unemployment spell is right censored. For any employment spell $m \in M$, $\tilde{w}_m$ is the observed wage corresponding to the $m-th$ job, t_m is the length of the $m-$th consecutive job, r_m is whether the $m-th$ job spell is right censored. q_m is the indicator whether the $m-th$ job is dissolved by the end of the job spell. Therefore, $q_m = 1$ when individual ends the $m-th$ job spell to be unemployed, $q_m = 0$ when the individual ends the $m-th$ job spell with another new job. We will firstly describe the likelihood function for a single unemployment/job spell. We then describe the likelihood function of an employment cycle as a combination of unemployment spells and job spells.

The likelihood contribution of an unemployment spell We first describe the likelihood contribution of an unemployment spell. The hazard rate is assumed to be

$$h_U = \lambda_U \bar{G}(\theta^*)$$

and the the density of the unemployment spell duration is

$$f_U(t_U) = h_U \exp(-h_U t_U)$$

The exact likelihood value from this unemployment spell would also depend on the censorship of the unemployment spell. When the unemployment spell is censored, then

$$l_U(t_U, r_U = 1) = \exp(-h_U t_U)$$

if the unemployment spell is completed, the exact likelihood value would be

$$l_U(t_U, r_U = 0) = h_U \exp(-h_U t_U)$$

The likelihood contribution from a job spell m We then describe the contribution of a job spell m to the likelihood function. The path dependence in our model is captured by the wage from last job (or the reservation wage w^* if the last spell is an unemployment spell). Given the wage w_{m-1} from the last spell, the distribution of the wage in any immediately successive spell is $\frac{f(w)}{\bar{F}(w_{m-1})}, w > w_{m-1}$. Given a random wage draw w_m from this distribution, the worker will only leave the current job spell for two reasons: (1) the current job may exogenously dissolve with rate η and the worker becomes unemployed. ($q_m = 1$) (2) the worker may move to an alternate firm with a better wage offer $w' > w_m$. ($q_m = 0$) Therefore, the total "total hazard" rate associated with this job spell is simply the sum of these two cases:

$$h_E(w_m) = \lambda_E \bar{F}(w_m) + \eta$$

The exact likelihood value given the true wage w_m depends on the right censorship r_u of the employment spell. It also depends on the reason why the job spell ends if it is not right censored. In summary, its likelihood value given the true wage w_m is

$$l_E(w_m, t_m, r_m, q_m|w_{m-1}) = \exp\left(-h_E(w_m)t_m\right)\left[\left(\lambda_E \bar{F}(w_m)\right)^{1-q_m}\eta^{q_m}\right]^{1-r_m}\frac{f(w_m)}{\bar{F}(w_{m-1})}$$

Now, one additional complication is that we do not know the true wage w_m but rather observe a noisy measure $\tilde{w}_m = w_m\epsilon$. Therefore, our likelihood value based on observed $\tilde{w}_m$ needs to integrate out the true wage w_m using its density function $w_m \sim \frac{f(w)}{\bar{F}(w_{m-1})}, w > w_{m-1}$

$$\begin{gathered} l_E(\tilde{w}_m, t_m, r_m, q_m|w_{m-1}) = \int_{w_{m-1}} \exp\left(-h_E(w_m)t_m\right) \\ \left[\left(\lambda_E \bar{F}(w_m)\right)^{1-q_m}\eta^{q_m}\right]^{1-r_m} \tfrac{1}{w_m} m\left(\tfrac{\tilde{w_m}}{w_m}\right)\tfrac{f(w_m)}{\bar{F}(w_{m-1})}dw_m \end{gathered}$$

where the term $\frac{1}{w_m}m\left(\frac{\tilde{w_m}}{w_m}\right)$ is the density of the observed wage $\tilde{w}_m$ under the log normal measurement error specification; the term $\frac{1}{w_m}$ is the Jacobian of the transformation

We now describe the likelihood function of a complete employment cycle. We focus on at most the two job spells ($M \leq 2$) in each employment cycle to reduce the computational burden. In the case when the employment cycle starts with an unemployment spell, we have

$$\begin{gathered} L_U(t_U, r_U, \tilde{w}_1, t_1, r_1, q_1, \tilde{w}_2, t_2, r_2, q_2) = \int_{w^*}\int_{w_1} h_U^{(1-r_U)}\exp(-h_U t_U) \\ \times\left\{\exp\left(-h_E(w_1)t_1\right)\left[\left(\lambda_E \bar{F}(w_1)\right)^{1-q_1}\eta^{q_1}\right]^{1-r_1}\tfrac{1}{w_1}m\left(\tfrac{\tilde{w_1}}{w_1}\right)\right\}^{1-r_U} \\ \times\left\{\exp\left(-h_E(w_2)t_2\right)\left[\left(\lambda_E \bar{F}(w_2)\right)^{1-q_2}\eta^{q_2}\right]^{1-r_2}\tfrac{1}{w_2}m\left(\tfrac{\tilde{w_2}}{w_2}\right)\right\}^{1-r_1} \\ \times\tfrac{f(w_1)}{\bar{F}(w^*)}\tfrac{f(w_2)}{\bar{F}(w_1)}dw_2 dw_1 \end{gathered}$$

In the case when the employment cycle starts with an employment spell, we have

$$\begin{gathered} L_E(\tilde{w}_1, t_1, r_1, q_1, \tilde{w}_2, t_2, r_2, q_2|w_1) = \\ \int_{w_1}\left\{\exp\left(-h_E(w_1)t_1\right)\left[\left(\lambda_E \bar{F}(w_1)\right)^{1-q_1}\eta^{q_1}\right]^{1-r_1}\right\}^{1-r_U} \\ \times\left\{\exp(-h_E(w_2)t_2)\left[\left(\lambda_E \bar{F}(w_2)\right)^{1-q_2}\eta^{q_2}\right]^{1-r_2}\tfrac{1}{w_2}m\left(\tfrac{\tilde{w_2}}{w_2}\right)\right\}^{1-r_1} \\ \tfrac{f(w_2)}{\bar{F}(w_1)}dw_2 \end{gathered}$$

Since we do not observe the precedent wage offers before the employment spell starts, we have no information about the density of the true wage w_1 at the beginning of the period.

Instead, we infer its distribution based on its measure $\tilde{w}_1$ from the relationship $w_1 = \tilde{w}_1/\epsilon$.

$$
\begin{gathered}
L_E(\tilde{w}_1, t_1, r_1, q_1, \tilde{w}_2, t_2, r_2, q_2) = \\
\int_{w^*}\int_{w_1}\left\{\exp\left(-h_E(w_1)t_1\right)\left[\left(\lambda_E\bar{F}(w_1)\right)^{1-q_1}\eta^{q_1}\right]^{1-r_1}\frac{\tilde{w}_1}{w_1^2}m\left(\frac{\tilde{w}_1}{w_1}\right)\right\} \\
\times\left\{\exp(-h_E(w_2)t_2)\left[\left(\lambda_E\bar{F}(w_2)\right)^{1-q_2}\eta^{q_2}\right]^{1-r_2}\frac{1}{w_2}m\left(\frac{\tilde{w}_2}{w_2}\right)\right\}^{1-r_1} \\
\frac{1}{\Gamma(\tilde{w}_1)}\frac{f(w_2)}{\bar{F}(w_1)}dw_2dw_1
\end{gathered}
$$

where $\frac{\tilde{w}_1}{w_1^2}m\left(\frac{\tilde{w}_1}{w_1}\right)\frac{1}{\Gamma(\tilde{w}_1)}, w_1 > w^*$ denotes the density of the "true" wage based on the its measured wage; the term $\frac{\tilde{w}_1}{w_1^2}$ is the Jacobian of the transformation, and $\Gamma(\tilde{w}_1)$ is a normalizing constant which ensures that the density integrates to unity.

A.2 Sample construction

A.2.1 Obtaining the dataset used in our analysis

This appendix describes the sample restrictions imposed to obtain the data subsample used for our analysis.

1. We restrict our sample to be individuals who are initially surveyed at 2013, with age between 25 to 60, with the number of individual 22,887.

2. We exclude samples whose marriage, education or gender information is not specified. The sample size is 20,643 after this step, reported in column 1.

3. We further drop the individuals with missing basic characteristics. This step leaves a sample size of 7,934 reported in column 2. (The main reason for the reduction in sample size was because cognitive ability was only measured in 2016)

4. We exclude job spells in which the labor force status transitions involve non-working states (any states other than full time, short time, part time, mini jobs or unemployment). This means individuals in our sample are ones who stay in the labor force at least once after 2013. We further exclude employment spells if the hourly wage associated with the first job in the employment spell is missing. This leaves our sample size to be 6,682, reported in column 3.

We can compare three samples: the raw sample (column 1), the sample with completed information (column 2), the working population (the final sample) for our estimation (column 3) As seen in the table below, the average characteristics are very similar across the three samples. The working population has somewhat higher education levels.

Table A1: Comparison of average characteristics for different subsamples†

	Raw sample	Completed information	Working population (Final sample)
Female	0.55	0.55	0.51
	(0.50)	(0.50)	(0.50)
	20643	7934	6683
Age	43.71	43.70	43.78
	(9.98)	(10.24)	(10.11)
	20643	7934	6683
Education	12.39	12.39	12.52
	(2.81)	(2.82)	(2.84)
	20643	7934	6683
Marriage	0.64	0.64	0.63
	(0.48)	(0.48)	(0.48)
	20643	7934	6683
Number of Children	1.04	0.97	0.94
	(1.19)	(1.16)	(1.12)
	20643	7934	6683
Cognitive ability	3.30	3.30	3.32
	(0.91)	(0.91)	(0.90)
	8267	7934	6683
Openness to experience	4.60	4.64	4.64
	(1.08)	(1.08)	(1.07)
	18167	7934	6683
Conscientiousness	5.85	5.85	5.85
	(0.81)	(0.79)	(0.78)
	18167	7934	6683
Extraversion	4.94	4.99	4.99
	(1.06)	(1.02)	(1.01)
	18167	7934	6683
Agreeableness	5.37	5.40	5.38
	(0.86)	(0.84)	(0.84)
	18167	7934	6683
Emotional stability	4.22	4.28	4.33
	(1.13)	(1.11)	(1.09)
	18167	7934	6682

†Every first line shows the sample mean. Every second line (number in parentheses) shows the standard deviation. and standard deviation. Every third line shows the number of observations.

A.2.2 Personality trait questionnaire

The table below describes the 15-item short version of the Big Five Inventory used in the GSOEP

Table A2: The 15-item short version of the Big Five Inventory in the GSOEP

I see myself as someone who ...	
Openness:	... is original, comes up with new ideas (+)
	... has an active imagination (+)
	... values artistic experiences (+)
Conscientiousness:	... does things effectively and efficiently (+)
	... tends to be lazy (-)
	... is relaxed, handles stress well (-)
Extraversion:	.. is communicative, talkative (+)
	... is outgoing, sociable (+)
	... is reserved (-)
Agreeableness:	... is considerate and kind to others (+)
	... is sometimes somewhat rude to others (-)
	... does a thorough job (+)
Neuroticism:	... gets nervous easily (+)
	... worries a lot (+)
	... is relaxed, handles stress well (-)

Note: (+) positively related with the trait; (-) negatively related with the trait.

A.3 Details Regarding Identification

We begin by considering identification of model parameters given access to the types of information available to us in the GSOEP dataset. This includes a continuous labor market history over a substantial period of time, in which the beginning and ending of job spells is available, as well as the beginning and ending dates of spells of nonemployment.[44] The job spells also consist of information on the wage earned, although this cannot be considered to be continuously measured. For the purpose of this discussion, we will consider model identification for the case in which the outside option used when determining the wage is the value of unemployed search. As was discussed above, in this case the wage is invariant over the course of a job spell, so that one observation of the wage at a job is sufficient for estimation purposes. Of course, when allowing for measurement error in wages, as we do, multiple measures of the wage at a job are valuable in terms of reducing measurement error. In the discussion below, though, we will assume that one wage observation is available for each job spell.

We will first discuss identification conditions for the model when there is no unobservable heterogeneity, that is, there exists a single value of z shared by all population members. This is the case often considered when structural search models are estimated, and relaxing this restriction is one of the contributions of this paper. The primitive parameters characterizing the model are time-invariant ability, a, and the distribution of match-specific productivity, θ, which has the parametric distribution $G(\theta|\Omega_\theta)$, with Ω_θ being a finite-dimensional parameter vector. In terms of Poisson event rate parameters, there are λ_U, λ_E, and η. In terms of what are essentially preference parameters, there is the discount rate ρ, and the flow utility parameter when unemployed, b. Finally, there is the surplus division parameter α which is the proportion of the match surplus received by the worker.

Our discussion begins by considering the case in which there is no measurement error and no on-the-job (OTJ) search, that is, $\lambda_E = 0$. Furthermore, we will assume that $\alpha = 1$, so that workers get the entire surplus from the employment match. In such a case, the wage received by an individual is given by $w = \alpha\theta = \theta$. This is the case considered by Flinn and Heckman (1982), hereafter referred to as FH, under the assumption that the data available were limited to CPS-like point-in-time information consisting only of the length of the in-progress unemployment spell if the individual was currently unemployed and the current wage if the individual was employed at the survey date.

Their principal findings were the following. Although the accepted wage distribution could be nonparametrically identified, there was no information in the data regarding re-

[44] In this paper we do not distinguish spells of unemployment from spells of nonparticipation, so that nonemployment is the state occupied by all individuals not currently employed.

jected offers, that is, any offers $w < w^*$, where w^* is the common reservation wage. Thus either one can assume that all offers are accepted (as in Flinn (2002)), in which case the support of the wage distribution must be restricted to a subset of $[w^*, \infty)$. Without this assumption, the accepted wage distribution is the wage offer distribution truncated from below at w^*, and the only way that the entire wage offer distribution can be estimated is under functional form assumptions.[45] In fact, FH show that even under parametric assumptions on $F(w)$, all parameters may not be identified unless the distribution F belongs to the class of distributions that they termed *recoverable.* Since it is natural to restrict the support of F to R_+, distributions of strictly non-negative random variables such as the truncated (from below) exponential and Pareto (with an unknown lower support point) were not recoverable, that is, the parameters characterizing these distributions could not be identified using only accepted wage information.

FH showed that the model, even with parametric assumptions on G (assuming that G was recoverable) did not allow point identification of the pair of parameters (ρ, b). Assuming no measurement error, they proposed an estimator that first estimated w^*, which is not a primitive parameter but instead characterizes the decision rule of the agents in the population. Their estimator was the (first) order statistic

$$\hat{w}^* = \min_{i \in S_E} \{w_i\},$$

where S_E is the set containing the indices of sample members who were employed at the time of the survey and w_i is the wage of individual i. Given the properties of the the distribution of order statistics, it is straightforward to demonstrate that $\operatorname{plim}_{N_E \to \infty} \hat{w}^* = w^*$, where N_E is the number of employed individuals in the sample. Moreover, since $\hat{w}^*$ converges at rate N_E to w^*, the asymptotic distribution $\sqrt{N_E}(\hat{w}^* - w^*)$ is degenerate with all of its mass at 0. Using this result, FH formed the concentrated log likelihood function which conditions on $\hat{w}^*$ to define maximum likelihood estimators of λ_U, η, and Ω_θ. Using the functional equation

$$w^* = b + \frac{\lambda_U}{\rho + \eta} \int_{w^*} (w - w^*) dG(\theta; \varpi),$$

FH replace the point-identified unknown parameters with consistent estimators,

$$\hat{w}^* = b + \frac{\hat{\lambda}_U}{\rho + \hat{\eta}} \int_{\hat{w}^*} (w - \hat{w}^*) dG(\theta; \hat{\varpi})$$

[45] If information on rejected wage offers is available, this could potentially be used to estimate the untruncated wage offer distribution. This information is rarely available, and even when it is its use in estimation is not entirely straightforward. This is due to the fact that some rejected wages are greater than accepted wages, which necessitates adding some type of measurement error process into the model.

to form the locus of points $\widehat{(\rho, b)}$ that solve the (estimated) functional equation. If an assumption is made regarding the value of either ρ or b, the other parameter is point identified.

In Flinn (2006), this basic model is extended to include Nash bargaining[46] over wages. In Flinn (2006) it was assumed that there was no OTJ search and in this case, under Nash bargaining, the wage is given by

$$w(\theta) = \alpha\theta + (1-\alpha)\theta^*, \tag{17}$$

where θ^* is the reservation match productivity value which is equal to the reservation wage (i.e., $\theta^* = w^*$), with

$$\theta^* = b + \frac{\alpha \times \lambda_U}{\rho + \eta} \int_{\theta^*} (\theta - \theta^*) dG(\theta; \varpi).$$

for the case in which $\lambda_E = 0$.[47] The wage function specified in (17) continues to apply in the model with OTJ search estimated in this paper (which assumes no Bertrand competition between firms) and assuming that the outside option utilized in setting wages is the value of unemployed search, V_U. The key thing to note about (17) is that it is linear in the random variable θ. Since

$$\theta = \frac{w - (1-\alpha)\theta^*}{\alpha},$$

the distribution of wages is given by

$$F(w) = G(\frac{w - (1-\alpha)\theta^*}{\alpha}),$$

with density

$$f(w) = \frac{1}{\alpha} g(\frac{w - (1-\alpha)\theta^*}{\alpha}).$$

The accepted wage distribution is truncated from below at θ^*, so that the this distribution is given by

$$F_A(w) = \frac{G(\frac{w-(1-\alpha)\theta^*}{\alpha}) - G(\theta^*)}{\tilde{G}(\theta^*)}, \ w \geq \theta^*$$

with density

$$f_A(w) = \frac{\frac{1}{\alpha} g(\frac{w-(1-\alpha)\theta^*}{\alpha})}{\tilde{G}(\theta^*)},$$

where $\tilde{G}(x) \equiv 1 - G(x)$. Flinn (2006) considered identification in the class of location-scale

[46]This is more accurately described as surplus division since model structure is not consistent with the axioms required for Nash's solution to the surplus division problem.

[47]For the case in which $\lambda_E > 0$, the expression for θ^* is given by equation 5.

distributions with support R_+. If G is a location-scale distribution, then

$$G(\theta; c, d) = G_0(\frac{\theta - c}{d}),\ \ \theta > 0,$$

where $c > 0$ is the location parameter and $d > 0$ is the scale parameter, with G_0 being a known distribution (that is, the distribution is known up to its location and scale). In this case, the accepted wage distribution is

$$\begin{aligned} f_A(w; c, d) &= \frac{\frac{1}{\alpha d} g_0\left(\frac{w-(1-\alpha)\theta^* - \alpha c}{\alpha d}\right)}{\tilde{G}_0\left(\frac{w^*-(1-\alpha)\theta^* - \alpha c}{\alpha d}\right)} \\ &= \frac{\frac{1}{d'} g_0\left(\frac{w-c'}{d'}\right)}{\tilde{G}_0\left(\frac{w^*-c'}{d'}\right)} \end{aligned}$$

which is the density associated with a random variable that has a truncated location-scale distribution with known G_0 and location parameter c' and scale parameter d', where

$$c' = (1-\alpha)\theta^* - \alpha c$$

and

$$d' = \alpha d.$$

Given access to a random sample of wages from the accepted wage distribution, $w_i,\ i = 1, ..., N_E$, and given a consistent estimator of w^*, $\hat{w}^*$, the (concentrated) log likelihood function for the sample is

$$\ln L(c', d' | \hat{w}^*) = -N_E \ln d' - N_E \ln \tilde{G}_0\left(\frac{w^* - c'}{d'}\right) + \sum_i \ln g_0\left(\frac{w_i - c'}{d'}\right),$$

and the maximum likelihood estimators of c' and d' are

$$\{\hat{c}', \hat{d}' | \hat{w}^*\} = \arg\max_{c', d'} \ln L(c', d' | \hat{w}^*).$$

These estimators are $\sqrt{N_E}$ consistent given the estimator of $\hat{w}^*$, but since $\hat{w}^*$ is an N_E consistent estimator of w^*, we have that

$$\underset{N_E \to \infty}{\text{plim}} \{\hat{c}', \hat{d}' | \hat{w}^*\} = \underset{N_E \to \infty}{\text{plim}} \{\hat{c}', \hat{d}' | w^*\},$$

that is, the location and scale parameter estimators using the concentrated log likelihood function have probability limits that are functions of the true parameter value w^*, not its estimator.

Proposition 1 *The parameter α is not identified if G is a location-scale distribution with unknown values of c and d.*

Proof. From the discussion above, if G is a (lower-truncated) location-scale distribution, we can obtain consistent maximum likelihood estimators of c' and d' (Pitman (1939)). Since $\theta^* = w^*$, and we have a strongly consistent estimator of w^*, we have a strongly consistent estimator of θ^*. The location and scale parameters c' and d' are functions of α, c, d, and θ^*, and the three unknown parameters are α, c, and d, given our consistent estimator of θ^*. We have

$$\begin{aligned} \widehat{c'} &= (1-\alpha)\hat{\theta}^* - \alpha c \\ \widehat{d'} &= \alpha d, \end{aligned}$$

which is a system of two nonlinear equations in three unknowns, α, c, and d. If c is known and d is unknown, then

$$\hat{\alpha} = \frac{\widehat{c'} - \hat{\theta}^*}{\hat{\theta}^* + c},$$

so that we have a consistent estimator of α. If c is unknown and d is known, then we have

$$\hat{\alpha} = \frac{\widehat{d'}}{d},$$

which is a consistent estimator of α. Thus a necessary condition for identification is that G is not a location-scale distribution with both location and scale parameters unknown. ■

As in the current paper, the assumption is often made that the distribution of θ is lognormal. The lognormal distribution is not a location-scale distribution, but $\ln\theta$ does have a location-scale distribution since it follows a normal distribution. We show now that α is identified under this functional form assumption when there exists a random sample of wages from the accepted wage distribution.

If θ has a lognormal distribution, then

$$G(\theta; \mu_\theta, \sigma_\theta) = \Phi\left(\frac{\ln\theta - \mu_\theta}{\sigma_\theta}\right),$$

where Φ denotes the c.d.f. of the standard normal, and where μ_θ is the mean of $\ln\theta$ and σ_θ is the standard deviation of $\ln\theta$. We will investigate identification under the lognormality assumption assuming that we have access to a random sample of N_E observations on accepted wages of individuals who entered the job spell from the unemployment state. In this case the (conditional, on employment) probability of observing an accepted wage less than or equal to w is given by

$$F_A(w) = \frac{G(\frac{w-(1-\alpha)\theta^*}{\alpha}) - G(\theta^*)}{1 - G(\theta^*).}$$

If G is lognormal, then we have

$$\begin{aligned} G(\frac{w-(1-\alpha)\theta^*}{\alpha}) &= \Phi\left(\frac{\ln\left(\frac{w-(1-\alpha)\theta^*}{\alpha}\right) - \mu_\theta}{\sigma_\theta}\right) \\ &= \Phi\left(\frac{\ln(w-(1-\alpha)\theta^*) - \ln\alpha - \mu_\theta}{\sigma_\theta}\right), \end{aligned}$$

so that

$$F_A(w) = \frac{\Phi\left(\frac{\ln(w-(1-\alpha)\theta^*)-\ln\alpha-\mu_\theta}{\sigma_\theta}\right) - \Phi\left(\frac{\ln\theta^*-\mu_\theta}{\sigma_\theta}\right)}{1-\Phi\left(\frac{\ln\theta^*-\mu_\theta}{\sigma_\theta}\right)},$$

which has the density

$$f_A(w) = \frac{\{(w-(1-\alpha)\theta^*)\sigma_\theta\}^{-1}\phi\left(\frac{\ln(w-(1-\alpha)\theta^*)-\ln\alpha-\mu_\theta}{\sigma_\theta}\right)}{1-\Phi\left(\frac{\ln\theta^*-\mu_\theta}{\sigma_\theta}\right)}.$$

As in Flinn and Heckman (1982), if we assume that wages are not measured with error, at least after some trimming has been applied to delete outliers, a super-consistent estimator of θ^* $(= w^*)$ is given by

$$\hat{\theta}^* = \min_{i \in S_E}\{w_i\},$$

where the set S_E includes the indices of all of the employment members in the sample. We then can define the concentrated conditional log likelihood function of the sample as

$$\ln L(w|\hat{\theta}^*) = \sum_{i \in S_e} \ln f_A(w_i|\hat{\theta}^*).$$

For sample member i, their contribution to the log likelihood function is given by

$$\ln L(w_i|\hat{\theta}^*) = -\ln\sigma_\theta - \ln(w_i - (1-\alpha)\hat{\theta}^*) - \frac{1}{2}\ln(2\pi) - \frac{1}{2}q_i^2 - \ln\left(1-\Phi\left(\frac{\ln\hat{\theta}^*-\mu_\theta}{\sigma_\theta}\right)\right),$$

where

$$q_i \equiv \frac{\ln(w_i - (1-\alpha)\theta^*) - \ln\alpha - \mu_\theta}{\sigma_\theta}.$$

The conditional maximum likelihood estimator is defined by

$$(\hat{\mu}_\theta, \hat{\sigma}_\theta, \hat{\alpha}) = \arg\max_{\mu_\theta, \sigma_\theta, \alpha} \sum \ln L(w_i|\hat{\theta}^*),$$

where the three first order conditions are

$$\begin{aligned}
\frac{\partial L(\hat{\Omega})}{\partial \mu_\theta} &= 0 = \sum \hat{q}_i - N_E \times h\left(\frac{\ln\hat{\theta}^* - \hat{\mu}_\theta}{\hat{\sigma}_\theta}\right) \\
\frac{\partial L(\hat{\Omega})}{\partial \sigma_\theta} &= 0 = -N_E + \sum \hat{q}_i^2 - \frac{N_E}{\sigma_\theta} + N_E \times h\left(\frac{\ln\hat{\theta}^* - \hat{\mu}_\theta}{\hat{\sigma}_\theta}\right) \times \left(\frac{\ln\hat{\theta}^* - \hat{\mu}_\theta}{\hat{\sigma}_\theta}\right) \\
\frac{\partial L(\hat{\Omega})}{\partial \alpha} &= 0 = -N_E - \frac{1}{\sigma_\theta}\sum \hat{q}_i \times \left(1 - \frac{w_i - (1-\hat{\alpha})\hat{\theta}^*}{\alpha\hat{\theta}^*}\right) \\
&\Rightarrow\ 0 = -N_E + \frac{1}{\hat{\sigma}_\theta\hat{\alpha}\hat{\theta}^*}\sum \hat{q}_i \times (w_i - \hat{\theta}^*),
\end{aligned}$$

where

$$h(x) \equiv \frac{\phi(x)}{1-\Phi(x)}$$

is the hazard function associated with the standard normal distribution. From these expressions, we can see that all three of the parameters are identified asymptotically in the sense that the three first order conditions are linearly independent. The first FOC is a function only of $\sum \hat{q}_i$. The second FOC is a function of $\sum \hat{q}_i^2$. The third FOC, associated with α, is a function of $\sum \hat{q}_i$ and $\sum(\hat{q}_i \times w_i)$. For the case in which θ is normally distributed, the FOC associated with α is only a function of $\sum \hat{q}_i$, so that the FOCs associated with μ_θ and α are linearly dependent. In this case there is no unique solution to the three equation system. We knew this to be the case from the necessary condition in the proposition.

Of course, the fact that the bargaining power parameter α is theoretically identified from only the accepted wage distribution in the lognormal case does not mean that it can be precisely estimated with relatively small sample sizes, even under "ideal" conditions in which all of the model assumptions characterize the data generating process (DGP). We therefore perform a sensitivity check by varying each parameter around its estimated value, but keeping other parameters unchanged. Figure A1 shows that, in practice, the surplus division parameter α is precisely identified around its optimal values.

A.3.1 Adding Heterogeneity to the Model

In many cases, researchers utilizing a structural approach deal with heterogeneity in observables by defining separate classes of individuals and then estimating the model separately for each class, often with no restrictions on parameter values across the classes. In such a case, consistency of m.l. estimators requires the sample size in each class grow indefinitely large. If there are K classes, with N_k individuals in each class, and the parameters characterizing the model for the k^{th} group denoted ω_k, then the (identified) parameter estimates, $\hat{\omega}_k$, are only consistent when $N_k \to \infty$, $k = 1, 2, ..., K$.

This method requires grouping observations into specified "bins," with the number of bins limited in practice due to the requirement that the sample size in each bin be sufficiently large so as to justify asymptotic approximations in deriving the sampling distribution of the estimator. We have taken an opposite tact however, and the goal of our estimation method is to consistently estimate primitive parameters even when observed heterogeneity is treated as continuous with no arbitrary aggregation used in order to create bins with relatively large numbers of observations. We begin with a vector of observed characteristics z_i for individual i, where z_i is a $1 \times (M+1)$ vector, the first element of which is a 1 for all i, so that there are M actual covariates. An individual's type, z_i, has varying impacts on the primitive parameters characterizing the search environment, with the impact in terms of primitive parameter j being given by

$$l_j(z_i\gamma_j),$$

where γ_j is an $(M+1) \times 1$ vector of weights attached to the vector of observed heterogeneity components, and where $l_j(\cdot)$ is the link function associated with parameter j. The purpose of the link function is to map the scalar index $z_i\gamma_j$, which takes values on the entire real line, into the appropriate parameter space for primitive parameter j. For example, many of the primitive parameters take values on the positive real line R_+. In these cases, we have used $l_j = \exp(\cdot)$, which is a common choice. A parameter of particular interest is α, which takes values in $(0, 1)$. In this case we use $l_\alpha = \exp(\cdot)/(1 + \exp(\cdot))$, the logit transformation, which is another common choice.

By using this specification of the impact of observed heterogeneity on model parameters, we are freed from the curse of dimensionality associated with the discretization of continuously-varying individual characteristics into "bins." The cost of our specification is the restriction of the manner in which the primitive parameters of the model can vary with the characteristics vector z_i. The linear index specification is roughly analogous to that imposed regularly in the linear regression context. One key difference is the fact that the impact of a given characteristic $z_{i,m}$ on a primitive parameter j is not independent of the values of other characteristics $z_{i,r}$, $r \neq m$, whenever the link function l_j is nonlinear, which is the case for all of the parameters that we estimate.[48] As we show in Table A3, the marginal impacts of z_{ij} are heterogeneous across individuals and depend on their values of other characteristics z_{il}, $l \neq j$.[49]

It is important to emphasize that the way in which we introduce observable heterogeneity into the model nests the homogenous case considered above. This is the case since the vector z_i includes a 1 as the first element of the vector for all z_i (for notational transparency, we will refer to the first element of the vector γ_j as element 0, with the conditioning variables z_i being in positions $1, ..., M$). Since we include this as the first element, the first element in any parameter vector γ_j corresponds to an "intercept" term. Then by restricting $\gamma[1 : M] = 0_{1\times M}$ we have the homogeneous model. Define the matrix of observable characteristics of the N sample members by

$$Z_{N\times(M+1)} = \begin{bmatrix} z_1 \\ z_2 \\ \vdots \\ z_N \end{bmatrix}.$$

Proposition 2 *If the homogeneous model is identified, the heterogeneous characteristic model is identified if and only if*

$$rank(Z) = M + 1.$$

Proof. In the homogeneous case, the score vector is defined by

$$\begin{aligned} \frac{\partial \ln L}{\partial \omega} &= \sum_{i=1}^{N} \frac{\partial \ln L_i}{\partial \omega} \\ &= \left(\sum_{i=1}^{N} \frac{\partial \ln L_i}{\partial \omega_1} \; \sum_{i=1}^{N} \frac{\partial \ln L_i}{\partial \omega_2} \; \dots \; \sum_{i=1}^{N} \frac{\partial \ln L_i}{\partial \omega_K} \right)'. \end{aligned}$$

Identification of the homogeneous model implies that there is an unique vector of values $\hat{\omega}$

[48] In principle, the mean of the $\ln \theta$ distribution can assume values on the entire real line so that a link function is not necessary in this case. However, we have normalized the value of this parameter so that $E(\theta) = 1$ for purposes of identifying the parameter vector γ_a.

[49] It is also important to recognize that our substantive restriction on the index function is that it be linear in the parameters. Powers of the individual characteristics, interactions between them, etc., can all be easily accommodated in principle, although precise estimation and interpretation of the parameter estimates becomes a challenge. This is why we opted for introducing the covariates in the manner in which we did.

that solve the system of equations

$$\begin{bmatrix} \sum_{i=1}^{N} \frac{\partial \ln L_i(\hat{\omega})}{\partial \omega_1} \\ \sum_{i=1}^{N} \frac{\partial \ln L_i(\hat{\omega})}{\partial \omega_2} \\ \vdots \\ \sum_{i=1}^{N} \frac{\partial \ln L_i(\hat{\omega})}{\partial \omega_K} \end{bmatrix} = 0_{K\times 1},$$

and that have the property that $\text{plim}_{N\to\infty} \hat{\omega}_N = \omega_0$, with ω_0 denoting the true parameter values. The value of the primitive parameter ω_j for an individual with characteristics z_i is given by

$$\omega_{ij} = l_j(z_i\gamma_j),$$

where the link function l_j is monotone increasing and everywhere differentiable on R. For the homogeneous model, we have $z_i = 1\ \forall i$, so that for the j^{th} parameter we have $\omega_{ij} = \omega_j = l_j(\gamma_{j,0})$. Given identification of the homogeneous model, then by the invariance property of the m.l. estimator,

$$\hat{\gamma}_{j,0} = l_j^{-1}(\hat{\omega}_j),\ \ j = 1, ..., K.$$

Given consistency of the maximum likelihood estimator $\hat{\omega}_j$, $\hat{\gamma}_{j,0}$, $j = 1, ..., K$, is consistent as well.

In the general heterogeneous case, we define the $K \times N$ matrix $\Delta(\gamma, Z)$ as

$$\Delta(\gamma, Z) = \begin{bmatrix} \frac{\partial \ln L_1}{\partial \omega_1}\frac{\partial l_1(\hat{x}_{1,1})}{\partial x} & \frac{\partial \ln L_2}{\partial \omega_1}\frac{\partial l_1(\hat{x}_{1,2})}{\partial x} & \cdots & \frac{\partial \ln L_N}{\partial \omega_1}\frac{\partial l_1(\hat{x}_{1,N})}{\partial x} \\ \frac{\partial \ln L_1}{\partial \omega_2}\frac{\partial l_2(\hat{x}_{2,1})}{\partial x} & \frac{\partial \ln L_2}{\partial \omega_2}\frac{\partial l_2(\hat{x}_{2,2})}{\partial x} & \cdots & \frac{\partial \ln L_N}{\partial \omega_2}\frac{\partial l_2(\hat{x}_{2,N})}{\partial x} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial \ln L_1}{\partial \omega_K}\frac{\partial l_K(\hat{x}_{K,1})}{\partial x} & \frac{\partial \ln L_2}{\partial \omega_K}\frac{\partial l_K(\hat{x}_{K,2})}{\partial x} & \cdots & \frac{\partial \ln L_N}{\partial \omega_K}\frac{\partial l_K(\hat{x}_{K,N})}{\partial x} \end{bmatrix},$$

where x is the argument of the link function and $\hat{x}_{ji} \equiv z_i\hat{\gamma}_j$. The solution to the first order conditions associated with the maximum likelihood estimator is given by

$$\Delta(\hat{\gamma}, Z)Z = 0_{K\times(M+1)}.$$

In the homogeneous case, $M = 0$ and we have

$$\Delta(\hat{\gamma}, Z) \times 1_{N\times 1} = 0_{K\times 1},$$

and we have assumed that $\hat{\gamma}$ is unique for this case. Given that Z is of full column rank, the columns of the matrix

$$\Delta(\gamma, Z)Z$$

are also of full column rank, so that there exists a unique solution

$$\Delta(\hat{\gamma}, Z)Z = 0_{K\times(M+1)}$$

for the case of $M \geq 0$. It is obvious that if $\text{rank}(Z) < (M + 1)$ there is no unique solution for $\hat{\gamma}$. ■

Since the covariate matrix Z that we use in estimation is of full column rank, the maximum likelihood estimator for our model is consistent when durations of unemployment and job spells are measured without error, which is virtually always assumed,[50] and when wages are measured without error as well.

A.3.2 Measurement Error in Wages

It is clear that wages recorded in any survey are measured with error. In a well-known validation study of earnings, wages, and hours of work using the Panel Study of Income Dynamics (PSID) instrument, Bound et al. (1994) find that measurement error is not a major problem in terms of respondent reports of annual earnings, but measures of reported hourly compensation contain a much larger amount of measurement error, with the proportion of ln wage variation attributable to measurement error reaching 50 to 60 percent. This is likely to be an upperward-biased estimate due to some problems in defining a "true" hourly wage given the compensation methods of the firm whose employees participated in the study,[51] but it remains the case that the results point to the likelihood that measurement error is a significant component of the total variance in wages.[52]

The presence of measurement error is required for us to define a maximum likelihood estimator for at least two reasons, of which one applies even to the homogeneous case. This is the fact that, in the case when firms do not compete in a Bertrand manner over an already

[50] For an exception, see Romeo (2001). Measurement errors in the starting and/or ending dates of spells in event history data are propagated throughout the remainder of the labor market state occupancy process. This makes the measurement error process not i.i.d. and much more challenging to introduce into the estimation of the model.

[51] The problem was that the rates of pay were set for activities performed by the worker, and that the worker could be assigned to multiple tasks within a pay period. Therefore, even if the worker was aware of the rate of pay at each of the tasks they performed, they may have found it difficult to recall the amount of time that the devoted to each task. Ultimately the employee may have found it difficult to recall their hourly rate of pay because there wasn't one, strictly speaking.

[52] Bound and Krueger (1991) perform a validation study of yearly income data gathered in the March supplement of the Current Population Survey using as the "true" measure of earnings that reported to the Social Security Administration. They find that the annual earnings measure that is self-reported by respondents has a high level of agreement with Social Security earnings, with only 15 percent of the total variance in annual earnings. However, they impose a large number of sample inclusion restrictions in order to perform their analysis, so that this should be taken as a lower bound. It also applies only to annual earnings, which Bound and Krueger (1991) find to contain much less measurement error.

employed individual, the individual will only leave their current job if the alternative job is associated with a higher match productivity value, and hence a higher wage. Thus, the probability of a wage decrease in a job-to-job transition is 0, whereas in the data this event is often observed. By adding other characteristics of remuneration, such as employer-provided health insurance (Dey and Flinn (2008)), it is possible to generate a positive probability of a wage decline associated with a job-to-job move, just as when firms compete via Bertrand competition (Postel-Vinay and Robin (2002)), however these models also impose constraints on the data generating process that are violated in the data.[53]

The second reason that measurement error is required is due to the relatively flexible way in which observable heterogeneity is introduced into the model. When one or more covariates Z are continuous, then the probability that any two individuals in the sample have identical values of Z is zero. This means that the primitive parameters will differ for any two sample members i and j, since $Z_i \neq Z_j$ for all $i \neq j$. In this case, for any individual i there will be a unique value of θ_i^*, and, in general, it follows that $\theta_i^* \neq \theta_j^*$ for all $i \neq j$.[54] This fact makes it impossible to use the order statistic estimator of FH for θ_i^* since the number of observations of accepted wages when the sample member i is coming from the unemployment state is small, and the asymptotic arguments for consistency in the FH case applied to the homogenous model with a large number of cross-sectional observations. This means that the concentrated likelihood approach in FH cannot be applied in this case.

The reservation match value of individual i is given by $\theta_i^* = \theta^*(z_i, \gamma_{-a})$, where γ_{-a} is the vector of parameters in the linear index functions for the primitive parameters with the exception of those characterizing a_i. The reservation wage is given by

$$w_i^* = l_a(z_i\gamma_a)\theta^*(z_i, \gamma_{-a}).$$

Based on the model specification, z_i, and γ, the likelihood that sample member i will accept a wage $w < w_i^*$ when they are unemployed is 0, and this is reason that measurement error must be introduced when no explicit (and complex) restrictions are imposed on the parameter space to ensure that w_i^* is at least as large as any wage that sample member i accepts when they are unemployed.

[53]For example, in Dey and Flinn (2005) the probability of leaving a job with employer-provided health insurance to accept one without such insurance is zero, whereas such transitions are observed in the data. The Bertrand competition model (Dey and Flinn (2005),Postel-Vinay and Robin (2002)) implies that real wages over a job spell at an employer should never decrease, which is observed in the data.

[54]We qualify this claim since it is possible that even though $Z_i \neq Z_j$ for all $i \neq j$, which implies that the primitive parameters will be different for i and j, the combination of primitive parameters for each could produce the same value of the reservation match value, so that $\theta_i^* = \theta_j^*$. Some further technical conditions would need to be added to ensure that this was not the case. For the purposes of this discussion it suffices to say that this is an extremely unlikely event.

As is commonly done, we assume classical measurement error which is identically and independently distributed within and across individuals and job spells. In particular, we assume that wage j observed in the observed labor market history of individual i is given by

$$\tilde{w}_{ij} = w_{ij}\varepsilon_{ij},$$

where ε follows a lognormal distribution, so that the density of ε is given by

$$m(\varepsilon) = \phi\left(\frac{\ln(\varepsilon) - \mu_\varepsilon}{\sigma_\varepsilon}\right)/(\varepsilon\sigma_\varepsilon),$$

where ϕ denotes the standard normal density. We impose the restriction that $\mu_\varepsilon = -0.5\sigma_\varepsilon^2$, so that $E(\varepsilon) = \exp(\mu_\varepsilon + 0.5\sigma_\varepsilon^2) = \exp(0) = 1$, and

$$\begin{aligned} E\tilde{w}_{ij} &= w_{ij}E(\varepsilon_{ij}) \\ &= w_{ij}\ \forall(i,j). \end{aligned}$$

The primitive parameters of the model are included in the matrix Γ which contains all of the parameters of the index functions. In addition, we must estimate the standard deviation of measurement error σ_ε. We know that the measurement error variance can be consistently estimated solely by using multiple measures of the wage at the same job under our modeling implication that real wages are invariant over the job spell. With a consistent estimator of σ_ε, we can evaluate the log likelihood function at any trial vector $\tilde{\Gamma}$. In order to evaluate the log likelihood function at $\tilde{\Gamma}$, it is required to first compute the reservation wage for each individual sample member at $\tilde{\Gamma}$, with $w_i^*(\tilde{\Gamma}) = w^*(\tilde{\Gamma}; z_i)$. The measurement error variance estimate is only required to evaluate individual log likelihood contributions involving wages. Then the estimates of Γ are determined as described in A.3.1.

A.3.3 Uniqueness of the Maximum Likelihood Estimator

We have not proved that the log likelihood is globally concave in Γ, however varying starting values of Γ_0 in the nonlinear optimization algorithm led to convergence at the same parameter value estimates. The log likelihood of the simple model of Flinn and Heckman (1982) was globally concave after conditioning on a consistent estimator of the (common) reservation wage $\hat{w}^*$. Our estimation is slightly more complex since we add on the job search (that is, $\lambda_E > 0$) and the surplus division parameter α. Neither addition is likely to generate nonuniqueness in the estimates of Γ from the score vector. The additional complication is that we have introduced observable heterogeneity in a flexible way. However, each element

Figure A1: Parameter identification in the homogeneous case

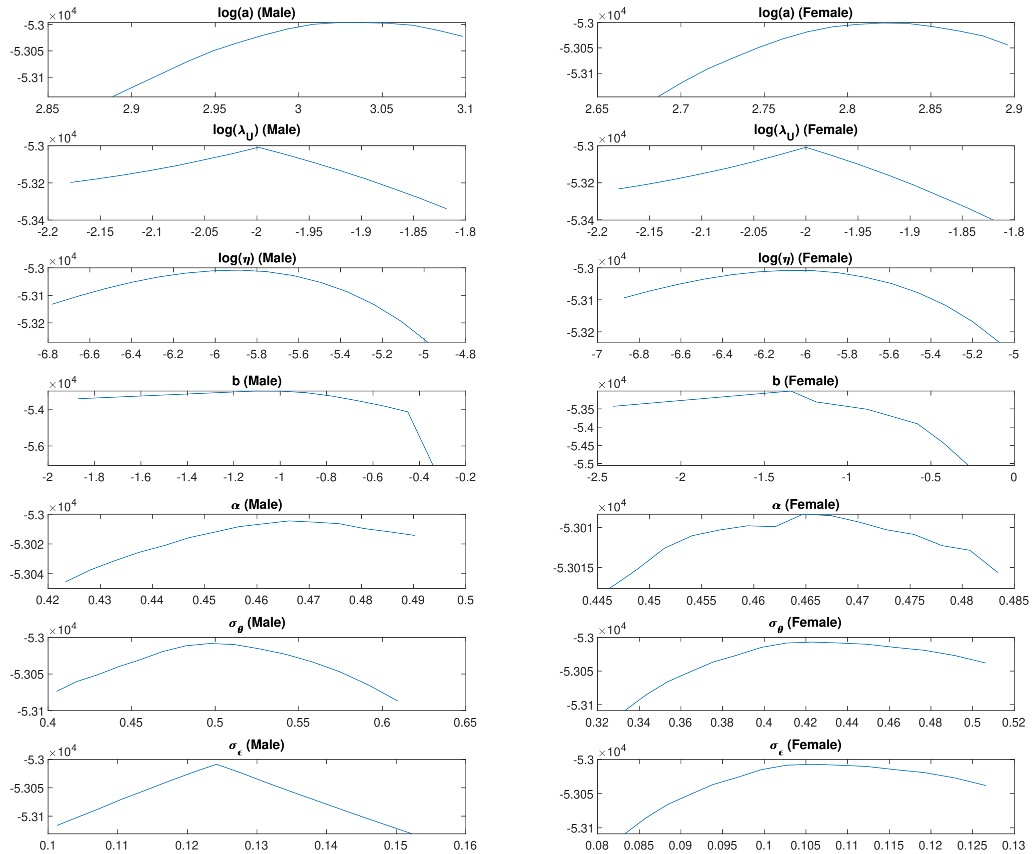

γ in the matrix Γ enters it's associated primitive parameter link function in a monotone way. Therefore, while we have not supplied a proof of uniqueness of the m.l.e., the finding of numerical stability with respect to changes in Γ_0 is not surprising.

Table A3: The marginal effect of personal characteristics (education, personality, cognitive ability) on job search parameters $\{a, \lambda, \eta, \alpha\}$

	Mean		25%		50%		75%	
	M	F	M	F	M	F	M	F
The marginal effect on ability da/dz								
Education	1.47	1.94	1.37	1.77	1.46	1.88	1.57	2.08
Cognitive ability	0.46	0.23	0.43	0.21	0.46	0.22	0.49	0.25
Openness to experience	0.16	0.17	0.15	0.15	0.16	0.16	0.17	0.18
Conscientiousness	0.20	0.47	0.19	0.43	0.20	0.45	0.22	0.50
Extraversion	0.37	0.35	0.34	0.32	0.36	0.34	0.39	0.37
Agreeableness	-0.21	-0.28	-0.22	-0.30	-0.20	-0.27	-0.19	-0.26
Emotional stability	0.37	0.21	0.35	0.19	0.37	0.20	0.40	0.22
The marginal effect on job arrival rate $d\lambda/dz$								
Education	0.048	0.040	0.032	0.028	0.045	0.037	0.057	0.049
Cognitive ability	0.012	0.031	0.008	0.022	0.011	0.028	0.014	0.038
Openness to experience	0.005	0.004	0.003	0.003	0.004	0.004	0.005	0.005
Conscientiousness	0.015	0.013	0.010	0.009	0.014	0.012	0.018	0.016
Extraversion	0.007	0.003	0.004	0.002	0.006	0.003	0.008	0.004
Agreeableness	-0.006	-0.013	-0.007	-0.016	-0.005	-0.012	-0.004	-0.009
Emotional stability	0.002	0.020	0.001	0.014	0.002	0.018	0.002	0.024
The marginal effect on job destructive rate $d\eta/dz$								
Education	-1.4E-03	-8.0E-04	-1.7E-03	-9.9E-04	-1.0E-03	-6.4E-04	-6.4E-04	-4.3E-04
Cognitive ability	3.8E-04	2.0E-04	1.7E-04	1.1E-04	2.8E-04	1.6E-04	4.7E-04	2.5E-04
Openness to experience	3.3E-04	-1.5E-04	1.5E-04	-1.9E-04	2.5E-04	-1.2E-04	4.1E-04	-8.3E-05
Conscientiousness	-7.0E-04	-5.8E-04	-8.6E-04	-7.2E-04	-5.2E-04	-4.6E-04	-3.2E-04	-3.1E-04
Extraversion	3.2E-04	1.3E-05	1.5E-04	7.2E-06	2.4E-04	1.1E-05	4.0E-04	1.7E-05
Agreeableness	2.8E-04	2.6E-04	1.3E-04	1.4E-04	2.1E-04	2.1E-04	3.5E-04	3.2E-04
Emotional stability	-4.0E-04	-2.1E-04	-4.9E-04	-2.7E-04	-3.0E-04	-1.7E-04	-1.8E-04	-1.1E-04
The marginal effect on bargaining power $d\alpha/dz$								
Education	0.033	0.027	0.033	0.027	0.033	0.027	0.033	0.028
Cognitive ability	0.003	0.002	0.003	0.002	0.003	0.002	0.003	0.002
Openness to experience	-0.014	-0.015	-0.014	-0.015	-0.014	-0.015	-0.014	-0.015
Conscientiousness	0.012	0.016	0.012	0.016	0.012	0.016	0.012	0.016
Extraversion	-0.011	-0.007	-0.011	-0.007	-0.011	-0.007	-0.011	-0.007
Agreeableness	-0.012	-0.026	-0.012	-0.026	-0.012	-0.026	-0.012	-0.025
Emotional stability	0.008	0.015	0.008	0.015	0.008	0.015	0.008	0.015