

Stockton, Isabel; Bergemann, Annette; Brunow, Stephan

Conference Paper

There And Back Again: Women's Marginal Commuting Costs

Beiträge zur Jahrestagung des Vereins für Socialpolitik 2016: Demographischer Wandel -
Session: Labor Markets: Gender Issues, No. B21-V1

Provided in Cooperation with:

Verein für Socialpolitik / German Economic Association

Suggested Citation: Stockton, Isabel; Bergemann, Annette; Brunow, Stephan (2016) : There And Back Again: Women's Marginal Commuting Costs, Beiträge zur Jahrestagung des Vereins für Socialpolitik 2016: Demographischer Wandel - Session: Labor Markets: Gender Issues, No. B21-V1, ZBW - Deutsche Zentralbibliothek für Wirtschaftswissenschaften, Leibniz-Informationszentrum Wirtschaft, Kiel und Hamburg

This Version is available at:

<https://hdl.handle.net/10419/145919>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

There And Back Again: Women's Marginal Commuting Costs*

January 2016

PRELIMINARY - PLEASE DO NOT CITE

Abstract: We use a partial-equilibrium model of job search with non-wage job attributes to estimate female workers' marginal willingness to pay to reduce commuting distance in Germany. This issue not only has implications for congestion policy, spatial planning and transport infrastructure provision, but is also relevant to our understanding of gender differences in labour market biographies in a more general sense. Gender differences in willingness to pay for job attributes could contribute significantly to observed disparities in a number of labour market behaviours and outcomes, such as participation, labour supply, wages and job mobility. Our analysis makes particular reference to heterogeneity by regional structure and local labour market conditions, and outlines ways of incorporating other types of heterogeneity into the analysis. Using a Cox model on a large administrative dataset, we find a marginal willingness to pay of €0.15 of women workers to reduce commuting distance by one kilometre.

Keywords: Commuting, marginal willingness to pay for job attributes, on-the-job search, Cox relative risk model, partial likelihood estimation, gender and parenthood in job search models, heterogeneity in job mobility

JEL Codes: C 41, J 13, J 16, J 31, J 62

*

Better estimates of workers' willingness to pay to reduce commutes are an important ingredient for welfare evaluations. On the one hand, this applies to traditional policy measures, such as transport infrastructure improvements. On the other hand, it refers to innovations in workplace organisation implemented by firms (such as telecommuting), as well as public policy measures designed to stimulate them. The latter are often advocated as improving day-to-day life for workers with care responsibilities. Analyses of marginal willingness to pay, especially focusing on women who more commonly carry such responsibilities, can help evaluate these claims empirically. Since commuting is a fixed cost associated to labour supply, estimating the utility cost of commuting for women can also help understand labour market participation decisions, and evaluate policies designed to raise it.

Looking beyond commuting, the issue of willingness to pay for job attributes has received renewed attention in labour economics in recent years. New methods address some of the econometric and conceptual problems of older, hedonic estimates, which have been shown to deliver biased results for markets with search frictions [Hwang et al., 1992, Gronberg and Reed, 1994, Hwang et al., 1998]. Gender differences in willingness to pay for job attributes could contribute significantly to observed disparities in labour market behaviours and outcomes, such as participation, labour supply, wages and job mobility. In this paper, we use a large administrative dataset to estimate female workers' marginal willingness to pay to reduce commuting distance, with a focus on regional differences. We find a moderate marginal willingness to pay of €0.19 to reduce commuting distance by one kilometre.

The paper is structured as follows: The first section provides an overview of the relevant literature. Subsequently, we specify a partial-equilibrium model of job search with jobs characterised by a wage and a commuting distance. Particular reference is made to differences by gender and the impact of the regional labour market situation, and a tractable estimator of marginal willingness to pay to reduce commuting distance is derived, following van Ommeren et al. [2000]. We then present an estimate of marginal commuting cost using a Cox model on an administrative linked employer-employee dataset. Finally, we outline avenues for future research.

1 Literature Review

Search Labour Markets and Gender From the vast literature on labour markets where the future stream of job offers is uncertain (for reviews, see Mortensen [1986], Mortensen and Pissarides [1999] and Rogerson et al. [2005]), sequential, partial-

equilibrium job search models are relevant to our estimation procedure. The first generation of search models centred on unemployed workers and assumed that once the searching worker has accepted a job offer, she either keeps it forever, or loses it for reasons exogenous to the model. From data on such a process, marginal willingness to pay for non-binary¹ job attributes cannot be recovered. On-the-job search was an early and intuitive extension due to Burdett [1978]. As long as the problem is stationary (the search environment does not depend on the current job, nor on the elapsed duration of the spell), this preserves the tractability and many of the core predictions of the original model.

Differences in the search process of male and female workers have been variously formalised as differences in job offer arrival rates, in job destruction rates, or in parameters governing exits into non-participation. In Bowlus [1997] and Bowlus and Grogan [2009], parameter estimates such as job finding and job destruction rates exhibit significant differences across genders, these differences are heterogeneous – in several instances, they even change sign – across education levels. This highlights the importance of accounting for interactions of gender with other determinants of labour market behaviour, instead of treating women as a homogeneous group. Differences in the value of non-market time, in search methods [Kuhn and Skuterud, 2004], in the propensity to bargain over wages [Hall and Krueger, 2012] or in geographical mobility [Blackaby et al., 2005], segregated professional networks [Montgomery, 1991, Mencken and Winfield, 2000], and discrimination [Bowlus and Eckstein, 2002, Flabbi, 2010] have all been suggested as underlying economic reasons for these differences. In our model, men and women are assumed to have different instantaneous utility functions over wages and commuting distances, leading to differences in marginal willingness to pay.

Commuting in Labour and Urban/Regional Economics Out of the many individual and collective decisions potentially affecting observed commuting distance, labour economics has chosen a different focus from urban and regional economics. This has resulted in fundamentally different models and empirical methods being developed in each field. In labour economics, residential location is usually treated as exogenous, turning commuting distance into a job or match attribute. Conversely, in urban economics, commuting distance is studied as a result of residential location decisions, which is compensated in the housing market. Possible interactions between these two markets are a central challenge in the modelling of commuting in either field, since adequate data,

¹Some authors [e.g. Villanueva, 2007] discuss a *marginal* willingness to pay for binary job attributes. However, it is plausible to assume that utility would change infra-marginally with the presence of a binary job attribute and would consequently not be continuous, much less differentiable.

as well as tractable models are often unavailable. Correspondingly, interpretations of observed gender differences also differ between the two fields.

Early studies on compensatory wage differentials for job (dis-)amenities augmented Mincerian wage regressions on education, labour market experience and tenure with measures of different job attributes, in particular risk of injury or death [Brown, 1980, Viscusi, 1978]. In addition to the mixed evidence it delivered, the hedonic approach has a number of theoretical weaknesses. Search frictions and unobserved productivity differences between workers and firms are not taken into account, leading to a substantial downward bias of OLS estimation of worker's marginal willingness to pay [Hwang et al., 1992, 1998]. One way to reduce bias from unobserved worker and firm heterogeneity is by instrumenting for non-wage job characteristics [e.g. Olson, 2002, for health insurance coverage]. However, for many job amenities including commuting distance, credibly exogenous variation in job attributes is difficult to find.

Concerns about endogeneity in hedonic estimates, as well as the difficulty of finding valid instruments for job attributes, gave momentum to the development of alternative estimation techniques. These methods put quit data and job spell durations in the focus. They improve earlier models of wage premia in a search environment [Khandker, 1988] by enabling the analysis of continuous job attributes, requiring fewer structural assumptions, and addressing issues such as sensitivity to outliers and measurement error [Eckstein and van den Berg, 2007]. The central idea is that more desirable job attributes decrease the probability of quitting in a search model with multi-dimensional jobs [Clark, 2001].

The model in van Ommeren and Fosgerau [2009] is a combination of a search model with a short-run model of commuting time for a fixed distance nested within the residential and workplace location decisions. The authors estimate the marginal cost of an hour's additional commuting at about twice the hourly wage using a discrete-time framework. This is much higher than estimates from models which focus on time costs only, which leads the authors to conclude that monetary costs make up about half of the overall amount at the margin. According to their findings, there is no significant difference between men's and women's marginal commuting costs.

Another strand of the literature uses parametric or semi-parametric duration models in continuous time to estimate parameters from job spell data [e.g. van Ommeren et al., 2000, which we will discuss in more detail when presenting the theory underlying our own project]. This is the approach which is most closely related to our work, and the theoretic model, which identifies marginal willingness to pay from observed job mobility, is similar to the long-run model by van Ommeren and Fosgerau [2009] we just discussed.

Russo et al. [2012] allow for heterogeneous effects by gender and find that women’s marginal willingness to pay is significantly higher than men’s. A number of other studies have explicitly examined gender differences in willingness to pay for a number of job attributes including workplace safety, type of tasks, promotion opportunities and different work schedule arrangements [Reed and Dahlquist, 1994, DeLeire and Levy, 2004, Felfe, 2012]. They have used duration models and conditional logit models, and studied the effect of job attributes on job-to-job mobility as well as mobility between labour market states around childbirth. However, commuting distance as a job attribute in a search model has rarely been analysed with an explicit gender focus so far.

Models of Commuting in Urban and Regional Economics Traditional models of commuting in urban and regional economics consider a household with one worker, often quietly presumed to be male. A number of models have extended this framework to accommodate household decision-making [e.g. Black et al., 2014] and the empirical literature has analysed commuting behaviour to explain differences by gender, chiefly women’s shorter commutes, a highly persistent finding across time and space [for Germany, see Auspurg and Schönholzer, 2013].

In an early contribution White [1986] uses an OLS regression of commuting time on income and demographic variables, with a focus on gender differences. However, her approach is unable to account for a number of unobserved differences between households classified as male- and female-headed, and is vulnerable to simultaneity bias. Timothy and Wheaton [2001] showed that the one-dimensional wage gradient predicted by monocentric models [White, 1986, Alonso, 1960, 1964] does not fit the observed wage variation in US Census data well. To explain the observed multidimensional variation, they therefore present a multinodal model. The spatial focus of the model can explain interactions between housing and labour markets, which most models in labour economics ignore. On the other hand, workers and firms are assumed to be matched without costly search, a substantial limitation from the labour-economic point of view.

Rouwendal and Rietveld [1994] extend a structural search model, assuming that the arrival rate of job offers depends on a worker’s location in space. They find that the commuting distance travelled by employed partners of heads of households decreased with the number of underage children in the household. The overwhelming majority of reported heads of households were men and of partners were women.² Since low-income

²This evidence supports the *household responsibility hypothesis*, which explains differences in commuting time with the uneven division of non-market work between men and women and consequent

households were oversampled, this suggests that the contrary results in White [1986] were driven by the over-representation of single mothers in the women's sample, instead of income *per se* playing a decisive role.

Models of residential location choice Estimating marginal commuting cost from a labour market search model requires the assumption that residential location is exogenous. If women are more likely to be *tied movers* or *tied stayers* [Mincer, 1978], whose private gains and losses are overridden by the net family utility change of a residential location decision, these models provide a better characterisation of female than of male workers' behaviour. Sermons and Koppelman [2001] and Clark et al. [2003] analyse how residential location decisions depend on men's and women's commutes in different household structures. Both papers find that location decisions are more sensitive to women's than to men's commutes, which is consistent with the household responsibility hypothesis.

There are two main estimation strategies in the choice and transportation economics literatures which estimate the value of travel time directly, without modelling job or residential location choices. On the one hand, stated preference methods study hypothetical choices from survey data. The results exhibit quite striking differences, even when similar models are applied [e.g. Calfee et al., 2001, Small et al., 2005]. It could be the case that the treatment of unobserved heterogeneity drives the wide variation in estimates, in which case the problem would be econometric in nature. Another quite likely but uncomfortable possibility is that stated preference data could simply be an unreliable signal of underlying preferences [Hensher, 2004].

On the other hand, revealed preference methods analyse observed choices nested within the commuting decision, such as mode, route, or vehicle choices [e.g. Brownstone et al., 2003, Lam and Small, 2001, both finding relatively high marginal willingness to pay for commuting time of more than 70% of the wage rate]. They circumvent the problem of biased reporting of willingness to pay, but data with sufficient variation of alternatives is not readily available, and the interpretation often extrapolates far beyond the range of commuting times actually observed. Older studies using revealed preference data also relied on strong assumptions on the shape of the utility function, as van Ommeren et al. [2000] point out.

higher productivity of women's non-market time

2 A Model of Job Search with Commuting

In this section, we will outline an on-the-job search model extended to two-dimensional jobs, closely following van Ommeren et al. [2000]. Without information on rejected job offers, marginal willingness to pay for non-binary job attributes cannot be recovered from a search process from unemployment. Voluntary job-to-job transitions, on the other hand, do identify this parameter.

Consider an employed worker in a job with wage w and commuting distance x , who receives alternative job offers (w^*, x^*) drawn from a distribution $F(w^*, x^*)$ according to a Poisson process with arrival rate λ . In addition to voluntary job-to-job transitions, employment spells end for exogenous reasons at rate δ . The expected discounted stream of utility from accepting job offer (w, x) over the whole of the life course is

$$\begin{aligned} \rho R(w, x) = & u(w, x) + \lambda \int \int \max\{0, R(w^*, x^*) - R(w, x)\} dF(w^*, x^*) \\ & + \delta(U - R(w, x)) \end{aligned}$$

where ρ is a discount parameter and U is the expected present value of unemployment. Lifetime utility is thus composed of an instantaneous component, a continuation value in case of a job switch and another continuation value in case of exogenous job loss.

The optimal strategy, as in the one-dimensional job case treated by Mortensen [1986], is myopic. The reason for this is that lifetime utility R depends on (w, x) only through instantaneous utility $u(w, x)$ and there are no transaction costs. Intuitively, whereas in a model without on-the-job search, a worker may “hold out” for a better offer, workers here have nothing to lose in accepting a job offer. They will still have an equal chance of receiving a better offer on the job.

Therefore, the worker pursues a reservation utility strategy: She accepts all job offers which offer a higher instantaneous utility than her present job, since the future stream of job offers is not affected by the job currently held. Formally, the set of job offers that are acceptable (i.e., strictly preferred to the current job) is

$$\varsigma(w, x) = \{(w^*, x^*) | u(w^*, x^*) > u(w, x)\}$$

This search and decision process leads to the following specification for the hazard rate from a job (w, x) :

$$\theta(w, x) = \delta + \lambda \int_{\varsigma(w, x)} dF(w^*, x^*) = \delta + \lambda(1 - F_u(u(w, x))),$$

i.e. the rate of exit from a job is given by the rate of exogenous exits into unemployment, plus the product of the rate of arrival of alternative offers and the probability that the offer will induce the worker to switch jobs. The second expression follows by substituting the above characterisation for the set of acceptable job offers, with F_u denoting the c.d.f. of $u(w, x)$.

As stated before, lifetime utility in this model depends on the wage and the commuting distance only through instantaneous utility. Therefore, the partial derivative of the hazard rate with respect to the wage w can be expressed as

$$\frac{\partial \theta(w, x)}{\partial w} = \frac{\partial \theta(w, x)}{\partial u(w, x)} \frac{\partial u(w, x)}{\partial w}$$

Clearly, an analogous statement holds for the derivative with respect to the commuting distance x .

This, in turn, gives us the equality stated by Gronberg and Reed [1994]: the instantaneous marginal rate of substitution or marginal willingness to pay for a job attributes is equal to the ratio of the marginal derivatives of the hazard rate:

$$\frac{\frac{\partial \lambda(w, x)}{\partial x}}{\frac{\partial \lambda(w, x)}{\partial w}} = \frac{\frac{\partial \lambda(w, x)}{\partial u(w, x)} \frac{\partial u(w, x)}{\partial x}}{\frac{\partial \lambda(w, x)}{\partial u(w, x)} \frac{\partial u(w, x)}{\partial w}} = \frac{\frac{\partial u(w, x)}{\partial x}}{\frac{\partial u(w, x)}{\partial w}} \quad (1)$$

Regional Labour Market Conditions As an extension to their basic model, van Ommeren et al. [2000] discuss the inclusion of business cycle effects in the model. They would affect the rate of arrival of job offers λ and/or the distribution $F(w, x)$ from which wage offers are drawn. Realistically, not only macroeconomic conditions at the national level should affect these two structural parameters of job search, but also regional trends, which therefore enter into the hazard rate.

In our empirical specification, we therefore include dummies for county settlement structure, as well as two indices of regional labour market conditions. These indices count regular employment relations in the individual's county of residence and in neighbouring districts, in her occupational field, by gender. The intuition is that the individual is likely to receive offers to work in her own profession, as well as in other professions within the same occupational field, which are defined with respect to similarity of tasks performed and skills required.

Importantly, we assume residential location to be fixed and define the local labour market environment around it. Therefore, the path of regional labour market devel-

opment facing an individual is the same, regardless of whether she changes jobs. This conserves the stationarity of the decision problem. A more complex model might treat regional labour market conditions as an additional job characteristic and consequently include incentives for relocation provided by regional differences in labour market conditions. But this would require modelling complex interaction between job and geographical mobility. Moreover, it would preclude the use of the simple estimator for marginal willingness to pay presented in van Ommeren et al. [2000].

Gender As explained in the literature review, a number of structural sources for different outcomes by gender in the search model have been suggested, such as different job offer arrival rates or different wage offer distributions. However, these would not affect marginal willingness to pay as a function.³ Since the hazard rate depends on (w, x) only through the instantaneous utility $u(w, x)$, the structural source of differences in marginal willingness to pay across the (w, x) -plain has to be differences in the instantaneous utility function.

As for the underlying causes of differences in instantaneous utility (which, in turn, generate differences in marginal willingness to pay), a common assumption in the literature is that women’s non-market time is more productive than men’s. This could be the case because they remain responsible for the bulk of household and child-rearing tasks. If social norms dictate that mothers should be nearby, for example to attend school-related meetings or events, or to be available in case of emergencies, the opportunity cost of commuting time would be much higher for women than for men and for mothers than for non-mothers. Other possible explanations include differences in access to a car, as suggested by Best and Lanzendorf [2005], or in the disutility from travelling, for example due to safety concerns.

3 Empirical Analysis

3.1 Data

The data set is a sample of the Institute for Employment Research’s *Integrated Employment Biographies*. Since the data comes from administrative social security records, they are more accurate than survey datasets commonly used in studies of commuting. For instance, they avoid problems of recall error in job spell durations and biased

³Marginal willingness to pay *at an observed wage* may differ if, as in most estimations in the literature as well as our baseline specification, it varies in the wage. In addition to purely empirical goodness-of-fit arguments, this specification can be justified by higher opportunity costs of commuting time for high-wage workers.

self-reporting of wages. The main sample combines a stock sample of all individuals employed on January 1st, 2000, with the inflow into employment after that date until December 31st, 2010, recording employment and unemployment spells in days. Self-employment is observed only in exceptional cases, since self-employed workers do not usually contribute to social security. The dataset is based on a 10% sample of all individuals with a Social Security number, going back to 1975. Employment spells in the sense of the model are times in regular full- or part-time employment with full social security contributions, which make up just over two-thirds of the full sample.

We construct an inflow sample from apprenticeship. This avoids the initial conditions problem as much as possible. If individuals are observed in several apprenticeships, we use the last one in chronological order. In order to minimise the impact of censoring whilst being able to exploit the multi-spell structure of the data [van den Berg, 2001], we only use the inflow in 2000 and 2001.

Apprenticeships and marginal employment (*geringfügige Beschäftigung* with either low total earnings or on short-term contracts, which are exempt from social security contributions) are excluded, since job mobility behaviour for these groups is likely to follow different patterns from regular workers and data is unavailable for earlier years. Moreover, we observe unemployment benefit claims, a number of other benefits such as early retirement programmes, and participation in active labour market policies such as publicly sponsored employment and training. Time out of the labour force appears as missing.

Spells of university graduates are excluded. There are two reasons for this: Firstly, university graduates are likely to be more geographically mobile, which makes the assumption of exogenous residential location more questionable. Secondly, wage data is top-coded, which could bias estimation, but rarely affects workers without a university degree.

Specification and Model Choice In the duration model of job mobility specified in this section, the failure event is a job ending for any reason. These include voluntary job-to-job transitions, layoffs and exits from the labour market. A job spell may be followed by a spell of missing data for several reasons: periods of full-time education, certain types of parental leave, time spent abroad and self-employment are all coded as missing. Voluntary transitions and layoffs are not unambiguously distinguishable from the data. But assuming that layoffs, like any exit events other than voluntary job-to-job transitions, are exogenous to the model, they will be captured by the baseline hazard. Since the estimator for marginal willingness to pay includes only partial derivatives of

the hazard, the layoff risk drops out.

This exogeneity assumption is questionable, since it is possible that, for example, the same set of reasons that induce a discontent worker to leave a job for another, may eventually induce him or her to exit the labour market entirely. However, the alternative model specification, where spells ending in any event other than a voluntary job-to-job transition are treated as censored, could be even more problematic. Censoring needs to be exogenous (non-informative) to the counting process of the event under study to allow consistent estimation. This assumption would be questionable for the exact same reasons as the exogenous layoff risk. In addition, some layoffs would probably be misclassified as voluntary transitions and vice-versa, exacerbating the bias.

A more complete model would directly model the hazards of leaving a job for other reasons in a competing risk framework. Identification of such a model is much more difficult than in the present case [van den Berg, 2001]. Bonhomme and Jolivet [2009] do specify a model where workers are at risk for different events ending a job and let the hazards depend on individual characteristics. They are not, however, allowed to depend on job characteristics, which they furthermore assume to be binary, a specification that would be of little interest in the case of commuting.

The wage data used in the estimation comes from social security records. Wage income above a certain threshold, which is usually adjusted annually, is not subject to social security contributions and therefore censored. To alleviate bias arising from this selection, we restrict our analysis to workers without a university degree, who less frequently earn wages above this cutoff values. Some systematic underreporting of higher education is known to occur. However, education information during job spells is considered most reliable, since it is employer-reported [Fitzenberger et al., 2005]. To minimise selection bias from underreported education, we smooth education, classifying individuals as university graduates after the first reporting of a university degree, even if the recorded variable switches to vocational training or no training afterwards.

The Cox Model The Cox relative risk model, introduced in Cox [1972], is a workhorse of regression-type duration analysis. It is a semiparametric model which leaves the baseline hazard unspecified and is estimated using a partial likelihood approach [Cox, 1972]. It is sometimes also referred to as the Cox proportional hazards model [e.g., Therneau and Grambsch [2000]]. However, the property of proportional hazards is not necessary for a Cox model. For example, Cox models may include time-varying covariates.

As Lancaster [1990, Chapter 9. 2] points out, economic theory rarely provides a sound base for predicting or explaining concrete survival times. This is in contrast

to medical or industrial applications, where there are often good reasons to assume an increasing or decreasing baseline hazard. In the absence of such guidance, the functional form assumptions required by parametric models are arbitrary. The Cox model, on the other hand, only uses information on the order in which individuals fail to estimate the effect of covariates, which is a terrain where economic models have much more to say.

In its most general formulation, the Cox model specifies the hazard as

$$\lambda(t, x) = \lambda_0(t)r(t, \mathbf{x})$$

for an individual with covariate vector \mathbf{x} .

Intuitively, the Cox model conducts a series of analyses at each observed failure time, where it asks the question: “Given that from the pool of survivors, one individual fails at this time, what is the probability of the failure occurring to the particular individual we actually observe to fail?”. However, it should be noted that a partial likelihood estimator which is used to estimate the Cox model cannot, in general, be interpreted as a (conditional) probability of any specific event. Instead, as the name suggests, it is a factor in a full likelihood and it can be shown that under certain non-restrictive conditions, it can be treated as equivalent to a likelihood for inference [Lancaster, 1990, p. 236ff and Kalbfleisch and Prentice, 2002, p. 100ff].

A simple example uses a proportional hazards specification of the form

$$\lambda(t|x) = \lambda(t) \exp(\mathbf{X}'\beta).$$

For the theoretical exposition, we abstract from the possibility of tied failures (which, in a continuous-time model such as this one, are a nuisance caused by imprecise measurement, not an actual feature of the process under study). Using the hazard rate as the central concept, the partial likelihood for the Cox model is the product over m terms, one for each observed failure time t_j . The j -th likelihood contribution is given by [Kalbfleisch and Prentice, 2002, p. 102]

$$L_j(\beta) = \frac{\lambda(t_j; x_j)dt_j}{\sum_{l \in R(t_j)} \lambda(t_j; x_l)dt_j}$$

where $R(t)$ is the risk set just before time t . Plugging in a relative risk model, it is easy to see that the baseline hazard cancels out. Thus, it remains unestimated and no functional form needs to be assumed. Instead of thinking in survival times and hazard rates, a survival model can be conceptualised as derived from a set of stochastic processes, namely a failure process N_i , an at-risk process Y_i and a covariate process X_i ,

which has to be predictable. This is the counting process paradigm used by Therneau and Grambsch [2000].

The switch from a survival time as an observation, to a left-open interval and an indicator of death, re-casts the central issue of censoring, which is no longer a problem of missing data, but a different state (similarly to the stochastic process approach of Lancaster [1990]). With this conceptualisation, the log partial likelihood (more intuitive than its exponentiated counterpart) is given by

$$l(\beta) = \sum_{i=1}^n \int_0^{\infty} \left[Y_i(t)r_i(t) - \log\left(\sum_j Y_j(t)r_j(t)\right) \right] dN_i(t)$$

where $r_i(t) = \exp[X_i(t)\beta]$ is the *risk score* or *relative hazard*.

Heuristically speaking, to estimate the effects of the covariates in a Cox model, the partial likelihood approach compares the expected value of the covariate process with its observed value. If the observed value is high relative to the expected one, this suggests a positive β [Kalbfleisch and Prentice, 2002, p. 103]. This idea is also the basis for the construction of Schoenfeld residuals, a popular test of the proportional-hazards assumption of the Cox model without time-variation.

As explained above, estimating a Cox model in continuous time means that ties arise only as a consequence of imprecise measurement. To handle them, we use the Breslow approximation [Breslow, 1974, Peto, 1972]. It calculates the partial likelihood assuming that both individuals recorded to fail at the same time are in the risk sets at each other's failure times. In the true model, of course, one individual must have failed first, so it is not part of the risk set of the individual who failed later. This approximation introduces a bias of the coefficients towards zero, but it is the least computationally demanding and performs well if ties are not too frequent [Kalbfleisch and Prentice, 2002, p. 105].⁴

Time-varying Covariates One of the advantages of the survival analysis methodology over more traditional regression models is its ability to model the influence of the current value of a time-varying covariate on the hazard. However, for a stochastic covariate process to be valid in a survival model, it has to be *predictable*. Heuristically,

⁴Therneau and Grambsch [2000] advocate the Efron approximation for use with SAS and S-plus, arguing that the increase in computing demand is negligible. In Stata, however, this is not the case. Since we work with a large dataset measuring durations with high precision, the simpler Breslow approximation is therefore preferable. The other alternatives, the exact marginal and exact partial methods which explicitly approximate all possible failure orders, are extremely demanding computationally. Moreover, Stata cannot automatically handle time-varying covariates with these methods, which means that records would need to be split to accommodate them. Together, these two aspects make memory demands of this strategy prohibitive for this application.

a predictable process does not look into the future: At each point in time, a realisation of a predictable process is determined only by information on the past history of the process itself and its covariates, not their future paths.⁵ Therneau and Grambsch [2000, p. 5ff] provide a good intuitive explanation in their introduction, relating the concept of predictability to a game of chance.

This condition can be problematic for the modelling of job mobility, if time-varying covariates are affected by agents' inside knowledge about the future path of the counting process. For example, a worker may learn about a job opportunity at a different firm and use that knowledge to negotiate his or her wage with her present employer. The employer may choose to raise the employee's wage to avoid losing her to a competitor. In this case, the covariate process – the wage – is affected by information on the future path of the job mobility process.

Therefore, its preimage is *not* contained in the filtration of the job mobility process up to the point where the wage negotiation takes place (it includes information on the counting process *after* this time), which makes the process non-predictable. In this case, the wage would not be a valid covariate in a model of job mobility. In our study, we only use workers without a university degree, which should help to reduce the impact of this problem somewhat, since individual wage negotiations are less common among this group.⁶

There are a few strategies that further research might pursue to evaluate the scale of the problem. For example, wage increases could be decomposed and only a predictable portion (such as average wage growth conditional on certain other covariates such as industry, region, educational attainment, etc.) used in the model. Another strategy would consist in studying a particular group of firms and workers, where individual wage negotiations are largely precluded, such as industries or firms with a high coverage of collective bargaining. In some institutional settings, the public sector also has wage schedules that could be modelled as a predictable process, but it is likely that job mobility will also be extremely low here, which make it a less appropriate case for study.

⁵For a more formal definition, recall that a stochastic process is a function that assigns random variables to time. It is predictable if and only if the preimage of this random variable is in the process's *filtration* [Kalbfleisch and Prentice, 2002, p. 157]. The filtration is usually defined as the sigma algebra generated by the history of the counting process of events up to and including the time t , as well as of the at-risk process and the covariate processes [Therneau and Grambsch, 2000, p. 18]. Left-continuity is a sufficient, but not a necessary condition for predictability.

⁶Hall and Krueger [2012] estimate a “dramatic” positive effect of education level on the probability of individual bargaining for the US (p. 64).

Functional Form of Covariates The aim of this section is to discuss alternative functional forms and decide on the most appropriate model. The standard Cox model assumes a linear form for the log relative risk, but a number of diagnostic tools are available to determine whether this simple specification fits the data well. There is a clear trade-off in model choice here: Linear and log-linear relative risk specifications are tractable and produce estimates of marginal willingness to pay to reduce commuting distance that are easy to interpret and to compare to previous work. However, they may oversimplify a complex relationship. Since the goal of this analysis is estimating a marginal cost of commuting, we prioritise finding a well-fitting specification for the effects of wage and commuting distance over other covariates.

We use fractional polynomials to find the best functional form for each covariate. This method runs through a pre-determined set of functions, and applies a formal deviance criterion to choose the best form. The available functions are degree-1 and -2 additive combinations of natural logarithms, fractional and integer powers (hence the name) from the set $\{-2, -1, -0.5, 0, 0.5, 1, 2, 3\}$. Evaluation of alternative specifications uses comparative measures based on the log partial likelihood, such that a higher-degree functional form is adopted if it leads to a significant change in the transformed likelihood. For a detailed discussion of fractional polynomials including an application to a Cox model, see Royston and Altman [1994].

The method of multivariate fractional polynomials expands this idea. An algorithm based on the fractional polynomial method just described runs through explanatory variables several times until convergence, in the order of increasing p-values. Covariates that have not been processed enter the model linearly, while those that the algorithm has already passed enter with the transformation that was judged optimal. To find the optimal functional form, the algorithm starts with the most complex (highest degree of freedom) transformation. If a reduction in degrees of freedom is not associated with a significant increase in deviance, the simpler functional form is adopted. Continuing this logic, variables that do not significantly reduce deviance vis-à-vis a model that excludes them are dropped.

We examine multivariate fractional polynomials for all continuous covariates. All of the variables are confirmed for inclusion in the model by the algorithm. The log-transformation of the wage, which is attractive on theoretical grounds and much-used in the literature, is confirmed. We use the functional forms suggested by this method for all controls, but let commuting distance enter the specification linearly to keep the estimate interpretable and comparable to previous studies.

Summary statistics are presented in table 1. The appendix gives details on variable

definitions, corrections applied and rules for inclusion in the sample.

Table 1: Summary Statistics by Gender

| | Women | | Men | |
|------------------------------|--------|-----------|--------|-----------|
| | Mean | Std. Dev. | Mean | Std. Dev. |
| Daily Wage | 55.34 | 23.40 | 66.81 | 25.95 |
| Commuting Distance | 12.15 | 12.99 | 13.62 | 14.01 |
| Job Duration (Years) | 2.82 | 3.34 | 2.76 | 3.48 |
| Unskilled Job | 0.04 | 0.20 | 0.03 | 0.17 |
| Workplace in Eastern Germany | 0.10 | 0.30 | 0.12 | 0.33 |
| Local LM: Men | 496.32 | 995.83 | 939.13 | 1,461.66 |
| Local LM: Women | 747.34 | 1,850.17 | 222.04 | 750.76 |
| Age | 26.55 | 7.11 | 23.29 | 5.43 |
| German nationality | 0.96 | 0.19 | 0.94 | 0.23 |

Wage and commuting distance statistics are averages over aggregate job spells to avoid arbitrary weighting.

Baseline Specification

As a baseline specification, we choose a Cox proportional hazards model with a log relative risk that is linear in commuting distance and uses the first-degree fractional polynomials for the other continuous covariates. This specification yields a marginal willingness to pay that depends linearly on the wage and is thus comparable to many previous estimates. Moreover, since time costs have previously been found to constitute a substantial part of overall commuting cost, it is theoretically plausible that they would vary in the opportunity cost of the worker’s time, i.e. the wage rate. The hazard rate can be expressed as

$$\lambda(t; x) = \lambda_0 \exp(\beta_w \ln(wage) + \beta_d distance + \beta_{\mathbf{x}} f(\mathbf{X})) \quad (2)$$

where the control vector \mathbf{X} includes age at the beginning of the spell, a dummy for unskilled occupation and sets of dummies for full-time work, regional categories of nationality, occupational category, county type and a proxy for the local labour market situation (for detailed variable definitions, see the appendix). Moreover, we include annual dummies to capture variation in the general macroeconomic conditions and national institutional environment, and regional GDP growth from the Federal Statistical Office [VGRdL, 2014].

Plugging the baseline functional form (2) into equation (1), marginal willingness to pay

Table 2: Baseline Specification

| | _t |
|----------------------------|----------------------------|
| | Baseline |
| Log Daily Wage | -0.664*** (0.013) |
| Commuting Distance | 0.00178*** (0.0000614) |
| Age | -0.0342*** (0.000936) |
| Unskilled Job | 0.175*** (0.026) |
| Full-Time | 0.211*** (0.0122) |
| Number of children (total) | -0.897*** (0.0275) |
| Childcare (worker proxy) | 0.00847 (0.00948) |
| Regional growth | -0.0000652 (0.00367) |
| Local LM: Men | 0.0303 (0.0104) |
| Local LM: Women | -0.00865 (0.0103) |
| Observations | 239388 |
| MWP | -0.00268 * <i>wage</i> |

Standard errors in parentheses

* $p < 0.0001$, ** $p < 0.00001$, *** $p < 0.000001$

is given by

$$MWP = \frac{\beta_d}{\beta_w \cdot \frac{1}{w}} = \frac{\beta_d}{\beta_w} w$$

The baseline estimation implies a marginal commuting cost of 14,8 cents per kilometer of commuting distance (i.e., half this figure per kilometre travelled) at the mean wage. This is substantially less than the €0,49 estimated for a mixed-gender sample of employees of Amsterdam’s VU university by Russo et al. [2012], and still somewhat less than the 0,4 Guilders or 18 Euro-cents estimated by van Ommeren et al. [2000] at the mean wage for men. This estimate, too, is for the Netherlands, so infrastructural, institutional and cultural differences might all contribute to a difference in marginal commuting cost. Also, it refers to the late 1980s, and it is plausible to assume that marginal commuting costs change over time, as prices for the monetary components of commuting cost, infrastructure, and the labour force composition change.

We fit separate models by regional type and find a lower marginal willingness to pay for workers in regions with agglomerations relative to the overall sample estimate. This is not consistent with the restriction to Amsterdam explaining the discrepancy. Education could be a factor driving the high estimate, since the Amsterdam sample contains a much higher proportion of university-educated workers than the one used by van Ommeren et al. [2000] and our sample excludes them altogether. However, although Russo et al. [2012] do not report separate estimates by education level, they find that faculty workers have a much lower marginal willingness to pay to reduce commuting distance than administrative staff, which makes this explanation unlikely. Another possibility is that the detailed controls used by Russo et al. [2012] eliminate some downward bias in the estimated commuting cost. However, some preliminary analyses using the method of Stratified Partial Likelihoods on our sample suggest that our baseline estimate is in fact biased upward by unobserved heterogeneity.

3.2 Separate Models by Regional Type

To analyse whether the search environment and, consequently, marginal commuting cost, differ according to the regional environment, we fit separate models by regional type. Table 3 shows the corresponding summary statistics and table 4 gives the estimation results.

Differences in observable characteristics mostly conform to expectations. Urban areas exhibit much greater employment opportunities, women earn higher wages and commute shorter distances there. The number of childbirths observed declines with the degree of agglomeration. This is an imperfect proxy of number of children, since we

Table 3: Summary Statistics by Region

| | | Urban | | Conurb. Features | | Rural | |
|-------------------------------------|----------|----------|-----------|------------------|-----------|--------|-----------|
| | | Mean | Std. Dev. | Mean | Std. Dev. | Mean | Std. Dev. |
| Daily Wage | | 57.69 | 24.54 | 53.69 | 22.20 | 50.98 | 20.97 |
| Commuting Distance | Dis- | 11.34 | 11.67 | 12.66 | 13.76 | 13.90 | 15.23 |
| Job Duration (Years) | Duration | 2.80 | 3.30 | 2.89 | 3.41 | 2.75 | 3.30 |
| Unskilled Job | | 0.04 | 0.19 | 0.04 | 0.20 | 0.04 | 0.19 |
| Workplace in Eastern Germany | East- | 0.07 | 0.26 | 0.11 | 0.32 | 0.19 | 0.39 |
| Any children present | | 0.48 | 0.50 | 0.51 | 0.50 | 0.52 | 0.50 |
| Maximum number of children observed | | 0.64 | 0.77 | 0.68 | 0.78 | 0.71 | 0.80 |
| Local LM: Men | | 827.94 | 1,306.80 | 185.26 | 257.89 | 112.83 | 167.95 |
| Local LM: Women | | 1,264.03 | 2,495.54 | 255.56 | 322.38 | 168.22 | 208.67 |
| Age in 2000 | | 24.12 | 6.82 | 23.79 | 6.89 | 22.71 | 6.48 |
| German nationality | | 0.95 | 0.21 | 0.97 | 0.16 | 0.99 | 0.11 |

Wage and commuting distance refer to averages over aggregate job spells to avoid arbitrary weighting.

might miss a second birth if a mother does not return to the labour market between two births. However, the share of women who are ever observed to have a child is also greater in more remote areas. It is worth noting that workers in areas with conurbational features are slightly more mobile overall, with a shorter average job duration. This is the only variable whose mean across different types of region reflects the pattern exhibited by marginal commuting cost, which is estimated to be highest in areas with conurbational features.

3.3 Baseline Model for the Men’s Sample

For comparison, we estimate the baseline model on a sample of men, constructed analogously to the women’s sample. Table 5 presents estimation results. Replicating a consistent result from the literature [e.g. Auspurg and Schönholzer, 2013], men are observed to commute greater distances in our sample (table 1).

Men’s estimated marginal willingness to pay is much smaller than for the whole women’s sample, and nearly all of the subsamples analysed. It should be kept in mind, however, that the assumption of exogenous residential location is more restrictive for the men’s sample. The assumption would be less appropriate than for the women’s sample if men are less likely to move for the sake of a partner’s career, or to stay despite

Table 4: Separate Models by Regional Type

| | Urban | Conurbational Features | Rural |
|----------------------------|--------------------------|--------------------------|--------------------------|
| | _t | _t | _t |
| Log Daily Wage | -0.711*** (0.018) | -0.606*** (0.022) | -0.651*** (0.038) |
| Commuting Distance | 0.00173*** (0.000088) | 0.00193*** (0.000105) | 0.00143*** (0.000145) |
| Age | -0.0329*** (0.00128) | -0.0350*** (0.00158) | -0.0364*** (0.00266) |
| Unskilled Job | 0.207*** (0.037) | 0.154 (0.043) | 0.0518 (0.0711) |
| Full-Time | 0.231*** (0.0172) | 0.182*** (0.0203) | 0.218*** (0.033) |
| Number of children (total) | -0.894*** (0.0387) | -0.872*** (0.0459) | -0.965*** (0.0725) |
| Childcare (worker proxy) | -0.0691*** (0.013) | 0.116*** (0.0203) | 0.111** (0.025) |
| Regional growth | -0.000921 (0.00504) | 0.00476 (0.00633) | -0.0155 (0.0117) |
| Local LM: Men | 0.0652 (0.018) | 0.008 (0.0163) | -0.00466 (0.0237) |
| Local LM: Women | 0.0245 (0.0173) | -0.0677* (0.0172) | -0.0338 (0.0246) |
| Observations | 119863 | 87139 | 32386 |
| MWP | -0.00243 * <i>wage</i> | -0.00319 * <i>wage</i> | -0.00219 * <i>wage</i> |

Standard errors in parentheses

* $p < 0.0001$, ** $p < 0.00001$, *** $p < 0.000001$

a good job offer in a different location, because of care responsibilities. Therefore, the substantial difference in marginal willingness to pay between genders may in part be due to compensation in the housing market for men. Men's estimated marginal commuting cost is only about 48% of women's, however, the implied total commuting cost at the mean wage and distance is 65% of women's, due to men's higher wages and longer commuting distances.

3.4 Separate models for mothers and non-mothers

We augmented the original dataset with information on employment interruptions due to childbirth, using a routine developed at IAB. We also constructed two proxies for childcare rationing, which could affect willingness to commute. The first of these proxies

Table 5: Baseline Specification: Men

| | | _t |
|--------------------------|-----------------------|-------------|
| Log Daily Wage | -1.017*** | (0.012) |
| Commuting Distance | 0.00132*** | (0.0000435) |
| Age -0.00843*** | (0.000812) | |
| Unskilled Job | 0.049 | (0.0221) |
| Full-Time | 0.250*** | (0.0191) |
| Childcare (worker proxy) | -0.00935 | (0.00734) |
| Regional growth | -0.00944 | (0.00303) |
| LM: Men | 0.0335* | (0.00786) |
| LM: Women -0.0306*** | (0.00538) | |
| Observations | 365051 | |
| MWP | -0.0013 * <i>wage</i> | |

Standard errors in parentheses

* $p < 0.0001$, ** $p < 0.00001$, *** $p < 0.000001$

is a local count of the number of childcare professionals, the second is the local childcare attendance rate for children under the age of three, which we were able to merge to our data for the years 2005-2010.

Table 6 provides summary statistics for the inflow sub-samples of women never observed to give birth, and for women before and after the birth of their first child. It should be kept in mind that this is an inflow sample following young women leaving an apprenticeship over a period of eight to ten years, so the sample of childless women is in fact, a sample of young childless women who may or may not have children later on.

The change to marginal commuting cost associated with childbirth is stark: As a fraction of the daily wage, mothers' marginal commuting cost after the birth of their first child is nearly double its previous values. This increase supports the hypothesis that women's higher commuting costs are to a large degree driven by the time cost of non-market childcare work. Since we derive the information on childbirth from social security records, and fathers are much less likely to interrupt their employment upon childbirth, an analogous analysis for men, which could shed additional light on gendered allocations of non-market work, is not possible with this dataset.

Both an increased marginal effect of commuting distance on the job mobility hazard and a decreased marginal effect of the wage contribute to the increase in women's marginal commuting cost upon childbirth, however, the proportional change in marginal effects is much larger for the wage. In the context of our structural model, this implies a substantially decreased marginal utility of wage, and a - less pronounced - increase in the disutility of commuting with childbirth. This is consistent with increased specialisation

Table 6: Summary Statistics for Mothers and Non-Mothers

| | | Childless women | | Before first birth | | After first birth | |
|------------------------------|----------|-----------------|-----------|--------------------|-----------|-------------------|-----------|
| | | Mean | Std. Dev. | Mean | Std. Dev. | Mean | Std. Dev. |
| Daily Wage | | 59.86 | 24.47 | 52.38 | 18.71 | 55.27 | 23.44 |
| Commuting Distance | Dis- | 12.71 | 13.09 | 12.14 | 13.18 | 12.13 | 12.98 |
| Job Duration (Years) | Duration | 3.15 | 3.62 | 1.87 | 2.00 | 2.83 | 3.35 |
| Unskilled Job | | 0.03 | 0.16 | 0.03 | 0.18 | 0.04 | 0.20 |
| Workplace in Eastern Germany | East- | 0.09 | 0.29 | 0.13 | 0.34 | 0.10 | 0.30 |
| Local LM: Men | | 528.20 | 982.15 | 476.46 | 985.01 | 495.80 | 996.46 |
| Local LM: Women | | 710.16 | 1,725.29 | 774.63 | 1,878.00 | 747.87 | 1,853.17 |
| Age in 2000 | | 21.94 | 5.46 | 20.63 | 2.87 | 23.94 | 6.89 |
| German nationality | | 0.97 | 0.17 | 0.96 | 0.19 | 0.96 | 0.19 |

Inflow sample from apprenticeship in 2000 and 2001. Ages refer to age in 2000.

Wages and commuting distances refer to aggregate job spell averages to avoid arbitrary weighting.

after the birth of the first child, with new mothers specialising in non-market work.

Mothers have higher marginal commuting costs than women before the birth of their first child and “forever” (during the observation period) childless women in all regional sub-samples. However, marginal commuting cost as shares of daily wages differ substantially across structurally different regions. Whilst marginal commuting for mother in rural areas are estimated to be higher than their counterparts’ without children within their area, they are lower than those of any group in areas with conurbational features.

To account for possible effects anticipating childbirth, I additionally estimated the model on a sample of women who are observed as mothers at some point, but only include observations at least three years before the birth of their first child. Marginal willingness to pay is similar for this sample, the entire pre-birth sample, and the sample of childless women. Interestingly, childless women’s job mobility is more sensitive both to commuting distance and to the wage than all three groups of mothers. This might be explained by other unmeasured job attributes, such as flexible hours or other family-friendly workplace policies, playing a greater role for mothers.

Interestingly, the childcare proxy is not significant. This result is robust to a number of specification changes, including interactions of childcare with wage and distance, the exclusion of controls for regional structure, and non-linear effects. The same happens for the other measure of childcare rationing, childcare attendance rates for children under the age of three. It is possible that these measures do not adequately capture the aspect

Table 7: Separate models for mothers and non-mothers

| | Before birth _t | 1st \geq 3 years be- fore 1st birth _t | After 1st birth _t | Childless women _t |
|-----------------------------|---------------------------|--|--------------------------|---------------------------|
| Log Daily Wage | -0.760*** (0.0252) | -0.734*** (0.0313) | -0.420*** (0.024) | -0.885*** (0.0211) |
| Commuting Distance | 0.00152*** (0.0000985) | 0.00169*** (0.000113) | 0.00169*** (0.000171) | 0.00196*** (0.0000824) |
| Age | 0.0163*** (0.00313) | 0.0133 (0.00378) | -0.0374*** (0.00177) | -0.00836*** (0.00148) |
| Unskilled Job | 0.143 (0.049) | 0.0954 (0.0625) | 0.201** (0.044) | 0.122 (0.0437) |
| Full-Time | 0.0323 (0.0215) | 0.0112 (0.0271) | 0.316*** (0.0205) | 0.106*** (0.0203) |
| Number of children | -1.240*** (0.116) | -0.853*** (0.14) | -1.981*** (0.136) | |
| Childcare (worker proxy) | -0.0147 (0.0152) | -0.00582 (0.0185) | 0.0108 (0.0215) | -0.0034 (0.0144) |
| Regional Growth | 0.00807 (0.00607) | 0.0092 (0.00827) | 0.00724 (0.00552) | -0.00454 |
| Local LM: Men | 0.0368 (0.0169) | 0.0348 (0.02) | 0.034 (0.0246) | 0.0429 (0.0162) |
| Local LM: Women | -0.0186 (0.0172) | -0.0189 (0.0211) | -0.0253 (0.0156) | -0.00753 |
| Observations | 65734 | 52047 | 55893 | 117761 |
| MWP | -0.002 * <i>w</i> | -0.00231 * <i>w</i> | -0.00403 * <i>w</i> | -0.00221 * <i>w</i> |

Standard errors in parentheses

* $p < 0.0001$, ** $p < 0.00001$, *** $p < 0.000001$

of rationing that affects commuting, which could plausibly be childcare institutions' opening hours. A complementary explanation is that childcare rationing might only affect the participation decision, not job mobility decisions taken once a mother is in the labour market.

Discussion

A potential weakness of the approach employed here is that selection into the labour force is unaccounted for. For some women with a high marginal willingness to pay to reduce commuting time, the offer distribution of wages and commuting distances may be insufficient to induce them to accept any job offer. This is because the high disutility they experience from commuting makes a job with the associated commute less attractive than non-employment. This could be especially relevant in a rural setting, where few jobs are available close-by. If sample selection thus excludes women with a particularly high marginal willingness to pay to reduce commuting distance, our estimates would be biased towards zero. This may partially explain the relatively low estimates. However, this sample selection problem affects any study based on job-to-job mobility, including Russo et al. [2012]. According to the OECD [2013], female labour force participation in Germany was up to five percentage points below that of the Netherlands in the first half of the 2000s, our observation window. In later years, participation rates were closer together, with 72% participation in Germany and 73% in the Netherlands at the end of the decade. Using data from Eurostat, Mischke and Wingerter [2012] even reports identical participation rates for 2011.

If differences in national participation rates are indeed an important contributor to the discrepancy in estimates, then it should diminish as participation rates converge. The average may, however, obscure important structural differences between the two countries. Regional culture and historical institutions in former East Germany continue to have an important impact on female labour force participation in Germany. All the Eastern *Länder* except Berlin exhibit higher female employment rates than either their Western counterparts, or any region of the Netherlands in most years [OECD, 2013]. Conversely, Dutch female labour force participation is highest in the especially densely populated Western region, which also has a higher GDP per capita than the other three regions [ibid]. Therefore, regional controls based on settlement structure or overall economic performance seem to be more closely correlated with participation in the Netherlands than in Germany. In particular, the results found by Russo et al. [2012] for metropolitan Amsterdam may not be comparable to our estimates for regions with

agglomerations, since sample selection bias affects both estimates in different ways.⁷

Another shortcoming of the dataset used in this paper is the lack of information on the the availability of telecommuting schemes. Firms which allow their employees to work from home more often effectively produce a better value of the job attribute “commuting time and expense”. However, the administrative dataset used here only records place of residence and place of work, not actual journeys required. A dataset with information on journeys made and/or journeys required by the employer, could reduce measurement error. Furthermore, it could provide a basis for an interesting analysis of firm-level choices on the production of job amenities. However, this data is unlikely to be available on a large scale, likely to suffer from measurement error and introduces new issues: if a telecommuting scheme is in place, but not widely used, is this because career penalties are associated with its use, or do employees trade off commuting cost with other job amenities, such as networking opportunities, an atmosphere more conducive to focused work, or a faster IT infrastructure?

Conclusion

Our estimates on a large administrative dataset indicate that women in Germany have a higher marginal willingness to pay to reduce commuting distance than men, which results in shorter observed commuting distances. Compared to previous estimates from search models, marginal willingness to pay is moderate. Moreover, there is evidence of heterogeneity by regional structure and age group. Workers in areas with conurbational features exhibit a higher marginal willingness to pay than their counterparts in either rural or urban areas. After the birth of their first child, women’s marginal commuting cost increases substantially.

A better understanding of how and why this pattern arises is crucial if valid recommendations for policy design in the areas of congestion policy, spatial planning and infrastructure are to be derived from the studies of marginal commuting cost. To achieve this in future research, we plan to better exploit the multi-spell structure of the data and analyse sub-samples of movers between areas of different structures and evaluate their marginal commuting cost in the pre- and post-relocation environments, as well as non-constant effects by commuting distance. This constitutes an important step to-

⁷For example, consider the following (stylised) scenario: Non-participating female workers in the Netherlands became discouraged because of a lack of available jobs, independent of commuting distance. Even a high effort generates a low arrival rate, but marginal willingness to pay for commuting is no different from the general population, which means that estimates are unbiased. Their counterparts in Germany, on the other hand, are deterred from entering the labour force because children’s daycare cannot accommodate a schedule extended by a long commute, generating a bias.

wards disentangling the influences of residential sorting, regional characteristics such as transport infrastructure, and/or differences in the distribution of commuting distances. This distinction is not generally made in existing labour-economic studies and will help bridge the gap to the commuting literature in urban and regional economics.

Going beyond observable differences between regions of residence, the multi-spell structure also makes the sample well-suited to the analysis of unobserved individual heterogeneity. This will be another focus of our project going forward, which has the potential to overcome important data-related limitations faced in previous work on this issue by addressing dynamic selection. To assess the effect of the changing composition of the group of employed women, we will conduct a number of additional analyses. Firstly, we will model an interaction effect between wages and county-year-specific childcare attendance rates. The important role of childcare rationing and cost for female participation is well documented in the literature [Connelly, 1992, Del Boca, 2002, Del Boca and Vuri, 2007]. In this model, wage is therefore assumed to have a baseline effect, referring to inframarginal participants, and an additional composition effect when more women enter into employment. We will cluster at the county level to account for the group structure thus introduced.

Furthermore, theoretical considerations suggest that selection into employment excludes workers with a low wage relative to the commuting distance. If marginal utility of wage, the source of heterogeneity, is distributed continuously or with more than two mass points, this means that the bias is more severe for low wage-commuting distance ratios. This hypothesis can be tested by including an interaction of wage and the inverse of commuting distance into the model. The resulting marginal willingness to pay will depend on both wages and commuting distances, making a direct presentation and interpretation of the estimate difficult. However, the direction of the effect could provide support for the theoretical argument that the baseline estimate is a lower bound.

Both the theoretical and the empirical literature give reason to suppose that women's shorter commutes are related to their greater share in non-market work. The large number of observations allows us to restrict the analysis to an inflow sample without overly great losses in precision. This enables us to account for unobserved heterogeneity, using the method of Stratified Partial Likelihoods [Ridder and Tunali, 1999]. This method identifies coefficients using within-individual variation and does not rely on independence of unobservables and covariates. Since unobservables that shape preferences for and productivity in market versus non-market work are likely to be correlated with covariates that arise from individual or household decision-making, such as wages, hours and education, this is an important advantage over random-effects (or in survival anal-

ysis terminology, shared-frailty) models. Whilst OLS estimates of wage premia have been augmented with fixed effects [e.g. Duncan and Holmlund, 1983, Villanueva, 2007], data limitations have prevented previous work using structural search models from using this within-person variation to effectively correct for bias from unobserved worker heterogeneity.

References

- William Alonso. A theory of the urban land market. *Papers in Regional Science*, 6(1): 149–157, 1960.
- William Alonso. *Location and land use. Toward a general theory of land rent*. Cambridge, Mass.: Harvard Univ. Pr., 1964.
- Katrin Auspurg and Thess Schönholzer. An Heim und Herd gebunden? Zum Einfluss von Pendelstrecken auf geschlechtsspezifische Lohnunterschiede. *Zeitschrift für Soziologie*, 42(2):138–156, 2013.
- Bundesinstitut für Bau-, Stadt- und Raumforschung BBSR. Kreise der Wohnungs- und Immobilienmarktbeobachtung. <http://www.bbsr.bund.de/BBSR/DE/WohnenImmobilien/Immobilienmarktbeobachtung/MethodenMaterialien/WIMKreise/WIMKreise.html?nn=446018>, 2006.
- Henning Best and Martin Lanzendorf. Division of labour and gender differences in metropolitan car use: an empirical study in Cologne, Germany. *Journal of Transport Geography*, 13(2):109–121, 2005.
- Dan A. Black, Natalia Kolesnikova, and Lowell J. Taylor. Why do so few women work in New York (and so many in Minneapolis)? Labor supply of married women across US cities. *Journal of Urban Economics*, 79:59–71, 2014.
- David Blackaby, Alison L. Booth, and Jeff Frank. Outside offers and the gender pay gap: Empirical evidence from the UK academic labour market. *The Economic Journal*, 115(501):F81–F107, 2005.
- Stephane Bonhomme and Gregory Jolivet. The pervasive absence of compensating differentials. *Journal of Applied Econometrics*, 24(5):763–795, 2009.
- Audra J. Bowlus. A search interpretation of male-female wage differentials. *Journal of Labor Economics*, 15(4):625–657, 1997.
- Audra J. Bowlus and Zvi Eckstein. Discrimination and skill differences in an equilibrium search model. *International Economic Review*, 43(4):1309–1345, 2002.
- Audra J. Bowlus and Louise Grogan. Gender wage differentials, job search, and part-time employment in the UK. *Oxford Economic Papers*, 61(2):275–303, 2009.
- Norman Breslow. Covariance analysis of censored survival data. *Biometrics*, 30(1):

89–99, 1974.

- Charles Brown. Equalizing differences in the labor market. *The Quarterly Journal of Economics*, 94(1):113–134, 1980.
- David Brownstone, Arindam Ghosh, Thomas F. Golob, Camilla Kazimi, and Dirk van Amelsfort. Drivers’ willingness-to-pay to reduce travel time: evidence from the San Diego I-15 congestion pricing project. *Transportation Research Part A: Policy and Practice*, 37(4):373–387, 2003.
- Bundesagentur für Arbeit. Klassifikation der Berufe 2010, Mai 2011. Nürnberg.
- Bundesanstalt für Arbeit. Klassifizierung der Berufe – Systematisches und alphabetisches Verzeichnis der Berufsbenennungen., September 1988. Nürnberg.
- Kenneth Burdett. A theory of employee job search and quit rates. *The American Economic Review*, 68(1):212–220, 1978.
- John Calfee, Clifford Winston, and Randolph Stempiski. Econometric issues in estimating consumer preferences from stated preference data: a case study of the value of automobile travel time. *Review of Economics and Statistics*, 83(4):699–707, 2001.
- Andrew E. Clark. What really matters in a job? Hedonic measurement using quit data. *Labour Economics*, 8(2):223–242, 2001.
- William A.V. Clark, Youqin Huang, and Suzanne Withers. Does commuting distance matter? Commuting tolerance and residential change. *Regional Science and Urban Economics*, 33(2):199–221, 2003.
- Rachel Connelly. The effect of child care costs on married women’s labor force participation. *The Review of Economics and Statistics*, pages 83–90, 1992.
- David R. Cox. Regression models and life-tables. *Journal of the Royal Statistical Society. Series B (Methodological)*, 34(2):pp. 187–220, 1972.
- Daniela Del Boca. The effect of child care and part time opportunities on participation and fertility decisions in italy. *Journal of Population Economics*, 15(3):549–573, 2002.
- Daniela Del Boca and Daniela Vuri. The mismatch between employment and child care in italy: the impact of rationing. *Journal of Population Economics*, 20(4):805–832, 2007.
- Thomas DeLeire and Helen Levy. Worker sorting and the risk of death on the job. *Journal of Labor Economics*, 22(4):925–953, 2004.
- Greg J. Duncan and Bertil Holmlund. Was Adam Smith right after all? Another test of the theory of compensating wage differentials. *Journal of Labor Economics*, 1(4): 366–379, 1983.
- Zvi Eckstein and Gerard J. van den Berg. Empirical labor search: A survey. *Journal of Econometrics*, 136(2):531–564, 2007.
- Christina Felfe. The willingness to pay for job amenities: Evidence from mothers’ return

- to work. *Industrial & Labor Relations Review*, 65(2):427–454, 2012.
- Bernd Fitzenberger, Aderonke Osikominu, and Robert Völter. Imputation rules to improve the education variable in the IAB employment subsample. FDZ Methodenreport, Nürnberg: IAB, 3 2005.
- Luca Flabbi. Gender discrimination estimation in a search model with matching and bargaining. *International Economic Review*, 51(3):745–783, 2010.
- Wilfried Görmar and Eleonore Irmen. Nichtadministrative Gebietsgliederungen und-kategorien für die Regionalstatistik. Die siedlungsstrukturelle Gebietstypisierung der BfLR. *Raumforschung und Raumordnung*, 49(6):387–394, 1991.
- Timothy J. Gronberg and William R. Reed. Estimating workers’ marginal willingness to pay for job attributes using duration data. *Journal of Human Resources*, 29(3): 911–931, 1994.
- Robert E. Hall and Alan B. Krueger. Evidence on the incidence of wage posting, wage bargaining, and on-the-job search. *American Economic Journal: Macroeconomics*, 4 (4):56–67, 2012.
- John C. Ham and Robert J. LaLonde. The effect of sample selection and initial conditions in duration models: Evidence from experimental data on training. *Econometrica*, 64(1):175–205, 1996.
- David A. Hensher. Identifying the influence of stated choice design dimensionality on willingness to pay for travel time savings. *Journal of Transport Economics and Policy*, 38(3):425–446, 2004.
- Hae-shin Hwang, William R. Reed, and Carlton Hubbard. Compensating wage differentials and unobserved productivity. *Journal of Political Economy*, 100(4):835–858, 1992.
- Hae-shin Hwang, Dale T. Mortensen, and William R. Reed. Hedonic wages and labor market search. *Journal of Labor Economics*, 16(4):815–847, 1998.
- John D. Kalbfleisch and Ross L. Prentice. *The statistical analysis of failure time data*. Wiley series in probability and statistics. J. Wiley, Hoboken, N.J., 2nd edition, 2002. ISBN 9781118032985.
- Rezaul K. Khandker. Offer heterogeneity in a two state model of sequential search. *The Review of Economics and Statistics*, 70(2):259–65, 1988.
- Peter Kuhn and Mikal Skuterud. Internet job search and unemployment durations. *American Economic Review*, 94(1):218–232, 2004.
- Terence C. Lam and Kenneth A. Small. The value of time and reliability: measurement from a value pricing experiment. *Transportation Research Part E: Logistics and Transportation Review*, 37(2):231–251, 2001.
- Tony Lancaster. *The econometric analysis of transition data*. Number 17 in Econometric Society Monographs. Cambridge University Press, 1990.

- F. Carson Mencken and Idee Winfield. Job search and sex segregation: Does sex of social contact matter? *Sex Roles*, 42(9):847–864, 2000.
- James Mincer. Family migration decisions. *Journal of Political Economy*, 87:749–773, 1978.
- Johanna Mischke and Christian Wingerter. Frauen und Männer auf dem Arbeitsmarkt: Deutschland und Europa. Technical report, Statistisches Bundesamt, 2012.
- James D. Montgomery. Social networks and labor-market outcomes: Toward an economic analysis. *The American Economic Review*, 81(5):1408–1418, 1991.
- Dale T. Mortensen. Job search and labor market analysis. *Handbook of labor economics*, 2:849–919, 1986.
- Dale T. Mortensen and Christopher A. Pissarides. New developments in models of search in the labor market. *Handbook of labor economics*, 3:2567–2627, 1999.
- OECD. Large regions, tl2: Regional labour market. *OECD Regional Statistics (database)*, 2013. doi: <http://dx.doi.org/10.1787/data-00528-en>. (Accessed on 11 June 2015).
- Craig A. Olson. Do workers accept lower wages in exchange for health benefits? *Journal of Labor Economics*, 20(S2):S91–S114, 2002.
- Richard Peto. Discussion of Professor Cox’s paper. *Journal of the Royal Statistical Society, Series B*, 34:205–207, 1972.
- William R. Reed and Julie Dahlquist. Do women prefer women’s work? *Applied economics*, 26(12):1133–1144, 1994.
- Geert Ridder and Insan Tunali. Stratified partial likelihood estimation. *Journal of Econometrics*, 92(2):193–232, 1999.
- Richard Rogerson, Robert Shimer, and Randall Wright. Search-theoretic models of the labor market: A survey. *Journal of Economic Literature*, 43:959–988, 2005.
- Jan Rouwendal and Piet Rietveld. Changes in commuting distances of dutch households. *Urban Studies*, 31(9):1545–1557, 1994.
- Patrick Royston and Douglas G. Altman. Regression using fractional polynomials of continuous covariates: parsimonious parametric modelling. *Applied Statistics*, 43(3): 429–467, 1994.
- Giovanni Russo, Jos van Ommeren, and Piet Rietveld. The university workers’ willingness to pay for commuting. *Transportation*, 39(6):1121–1132, 2012.
- M. William Sermons and Frank S. Koppelman. Representing the differences between female and male commute behavior in residential location choice models. *Journal of Transport Geography*, 9(2):101–110, 2001.
- Kenneth A. Small, Clifford Winston, and Jia Yan. Uncovering the distribution of motorists’ preferences for travel time and reliability. *Econometrica*, 73(4):1367–1382,

2005.

- Terry M. Therneau and Patricia M. Grambsch. *Modeling survival data: extending the Cox model*. Springer, New York, 2000.
- Darren Timothy and William C. Wheaton. Intra-urban wage variation, employment location, and commuting times. *Journal of Urban Economics*, 50(2):338–366, 2001.
- Gerard J. van den Berg. Duration models: specification, identification and multiple durations. *Handbook of econometrics*, 5:3381–3460, 2001.
- Jos van Ommeren and Mogens Fosgerau. Workers’ marginal costs of commuting. *Journal of Urban Economics*, 65(1):38–47, 2009.
- Jos van Ommeren, Gerard J. van den Berg, and Cees Gorter. Estimating the marginal willingness to pay for commuting. *Journal of Regional Science*, 40(3):541–563, 2000.
- Statistische Ämter des Bundes und der Länder VGRdL. Gross domestic product – at current prices – in Germany by Bundesländer. <http://www.vgrdl.de/VGRdL/tbls/tab.asp?lang=en-GB&rev=RV2014&tbl=tab01>, 2014.
- Ernesto Villanueva. Estimating compensating wage differentials using voluntary job changes: evidence from Germany. *Industrial & Labor Relations Review*, 60(4):544–561, 2007.
- W. Kip Viscusi. Wealth effects and earnings premiums for job hazards. *The Review of Economics and Statistics*, 60(3):408–416, 1978.
- Michelle J. White. Sex differences in urban commuting patterns. *American Economic Review*, 76(2):368–72, 1986.

4 Appendix

Variable Definitions

Occupational Characteristics The dataset contains two **occupation** variables, a 3-digit variable based on the 1988 classification [Bundesanstalt für Arbeit, 1988], and a 5-digit variable based on the most recent classification [Bundesagentur für Arbeit, 2011]. Since the observation window ends in 2010 and re-coding of older observations to the 2010 system is not error-free, the older variable is likely to be more accurate. However, the 2010 classification combines a horizontal (occupation) and a vertical (skill level) dimension. I recover the skill level information, which is absent from the older variable at the available aggregation level. For the horizontal occupation information, I match the 1988 information to 53 task-based *occupational fields*, as defined by the Institute for Vocational Education and Training.

Skill distinguishes four levels, with the following interpretation:

- *un- or semiskilled activity*: simple or routine tasks of little complexity, formal vocational training not usually required
- *professional activity*: requires sound professional knowledge and skills, two- to three-year vocational training
- *complex specialised activity*: requires specialised knowledge and skills, involves planning and management responsibility, requires a second vocational qualification (*Meisterprüfung*, master craftsman’s diploma), technical training or academic bachelor degree (EQF level 6)
- *highly complex activity*: requires expert knowledge, involves leadership and managerial responsibility, at least four years’ academic training

Regional Characteristics I match the individual data to classification which distinguishes 9 **types of district**, due to the Federal Institute for Research on Building, Urban Affairs and Spatial Development (BBSR). They are based on administrative districts, but differ from them where administrative divisions group structurally different areas into one unit. The resulting “analytical” districts are subdivisions of three types of region:

- *Regions with large agglomerations* are subdivided into core cities, densely populated areas, populated areas, and rural areas
- *Regions with conurbational features* are subdivided into core cities, populated areas, and rural areas
- *Rural regions* are subdivided into densely and sparsely populated rural areas

For more detailed information on the classification, see Görmar and Irmen [1991] or the institute’s online information [BBSR, 2006].

Local Labour Market Proxies Local labour market conditions could have a substantial impact on mobility behaviour if residential location is largely fixed. To capture this effect, I use proxies for local labour market conditions, based on a count of jobs by region and occupational field. Each spell is assigned the number of individuals employed full-time times the days the employment lasted, in her own county and in the neighbouring counties, in her occupational field. In the baseline specification, I disaggregate this measure by gender.

Treatment of Missings Individuals' **Nationality**, while not necessary constant, is unlikely to change back and forth, more so in a short time span. Therefore, up to two subsequent missing values were filled if coded nationality was identical for the individual before and after the gap arising due to the missing values. This affects a small part of the sample, since less than 1 percent of observations originally had nationality coded as missing.

If the **wage** was coded as zero or missing in a subspell, but valid wage information was available in another spell within the same aggregate job spell (a continuous employment at the same firm), the valid wage information was extended to the missing observation.

Children The variable **Number of children** is not time-varying, but refers to the total number of childbirths observed.

Sample Construction

Employment spells according to the above-mentioned definition are included in the sample if:

- they are not overlapped by a spell in registered unemployment, a programme to support the long-term unemployed, publicly sponsored employment, retirement, a seasonal or temporal work placement organised by the employment agency, or a mobility-related subsidy. Small overlaps of up to three days are tolerated. Individuals are coded as registered unemployed if they earn wages below the minimum subsistence level and are thus eligible for unemployment benefits, even though they are employed. This implies that the wage paid by the employer is not the wage actually perceived by the worker, who faces a wage distribution that is truncated at the minimum subsistence level (which, however, varies according to household composition).

A similar distortion of the wage-commuting trade-off arises in the case of a mobility subsidy. Participants in active labour market programmes, on the other hand, do not choose their place of work, and their behaviour can therefore not be adequately reflected in the model. Biases that may arise from selection into non-subsidised employment, or into employment more generally, are not addressed here. This excludes less than 3,3% of spells.

- the individual is never recorded as having a university degree, with certain corrections applied. To avoid complications arising from the decision to return to

education, we do not include employment spells before university graduation. We do not know if individuals acquire a university degree after the end of the observation window. However, eight to ten years after leaving vocational education, this probably does not apply to many individuals.

- they belong to a job identified as the main job at that time (for details, see the appendix). This leads to less than 0,6 % of spells being dropped.
- they are not part of a seasonal work pattern, i.e. the worker does not return to the same employer without an intervening spell at a different firm. Spells with the same employer with gaps of up to a week are considered part of a single job to avoid misinterpreting administrative delays as seasonal work. About a quarter of spells are dropped for this reason.
- they are part of the inflow sample between January 1st, 2000, and December 31st, 2001. The data is right-censored on December 31st, 2010.
- they last for more than 60 days. Location choice for temporary work is likely to take place under different conditions than for permanent employment. Moreover, most temporary contracts are already dropped as marginal employment. This excludes less than 2,8% of spells.
- the implied monthly wage is within the limits that make it non-exempt from social security contributions (*Geringfügigkeitsgrenze* and *Beitragsbemessungsgrenze*). Due to different timings of reports, wage information in some spells which are not actually subject to contributions was included in the original dataset.

These criteria are applied in the stated order, e.g., if short and long job spells form a seasonal work pattern, they are dropped, even if the short spells would later be excluded by the two-months rule.

Spells where either the place of work or the place of residence was missing or invalid, or where an individual was recorded as living in, or a firm recorded as being located in, two or more different zip code areas at once, were excluded since no valid commuting distance could be determined. The final sample of women contains just over seven million spells.

Short Spells Temporary workers whose contracts last less than two months are usually not liable to pay full social security contributions, which precludes their inclusion in the sample. Spells of under two months could be due to exceptions in the social

security treatment, early firings, miscoded part-time work, or misreported dates, which are difficult to disentangle. Moreover, the optimisation process underlying short-term job location may differ substantially from the one related to long-term job mobility decisions, and temporary residential relocations are likely to not appear in the data, which makes distance calculations unreliable. Therefore, spells of under two months are dropped.

Treatment of overlapping employment spells Overlapping spells present a challenge to the model, since neither the theory nor the software enables an agent to be in two states at once. Separate states for each combination (such as, “working while receiving benefits”, “working two jobs”, etc.) would have to be defined, and the concepts of willingness to pay for job attributes do not readily translate to such an environment. Therefore, I make the simplifying assumption that individuals have one main job, and mobility behaviour in any other jobs is not reflected in the model. Cases where no clear hierarchy of parallel jobs can be determined are excluded.

Multiple job spells with different employers at the same time Overlapping job spells of the same individual with different employers are excluded, except in the following cases:

- If the overlap is less than two weeks, both spells are included, with the transition assumed to occur at the start of the overlapping period.
- If one and only one of the jobs lasts for less than a year and the other one is at least three times as long, the longer spell is included.
- If one of the jobs is full-time whereas the other one is part-time, the full-time spell is included in the sample

The three criteria are hierarchical, i.e. the job duration criterion applies only if the two-week criterion does not, etc.

Multiple sub-spells with the same employer In this case, the outcome - job mobility - is unaffected, and the only question is which values of time-varying covariates are valid at which point in time. Pairs of these spells were split. The spell created from the overlapping time has the covariates of the two original spells if they are non-contradictory, otherwise, the covariate is set to missing. In the case of conflicting wage information, if the difference is less than 5%, the mean is used.

Browsing the data where spells overlap suggest that while some probably refer to changing wages, others appear to refer to bonuses instead, which would imply that the true wage is the sum of both recorded wages. Separating the two cases would most likely involve (more) arbitrary cutoffs or lists of professions, and less than 1 % of spells are affected, so no attempt at this as made. In the rare case of triple or greater multiple overlaps which only affects about 1 in 2000 spells, the overlapping portions were dropped without any corrections to the covariates.