

Lentz, Rasmus; Tranæs, Torben

Working Paper

Marriage, wealth, and unemployment duration: a gender asymmetry puzzle

IZA Discussion Papers, No. 1607

Provided in Cooperation with:

IZA – Institute of Labor Economics

Suggested Citation: Lentz, Rasmus; Tranæs, Torben (2005) : Marriage, wealth, and unemployment duration: a gender asymmetry puzzle, IZA Discussion Papers, No. 1607, Institute for the Study of Labor (IZA), Bonn

This Version is available at:

<https://hdl.handle.net/10419/33371>

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.

IZA DP No. 1607

Marriage, Wealth, and Unemployment Duration: A Gender Asymmetry Puzzle

Rasmus Lentz
Torben Tranæs

May 2005

Marriage, Wealth, and Unemployment Duration: A Gender Asymmetry Puzzle

Rasmus Lentz

Boston University

Torben Tranæs

*Rockwool Foundation Research Unit,
CESifo and IZA Bonn*

Discussion Paper No. 1607

May 2005

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0

Fax: +49-228-3894-180

Email: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of the institute. Research disseminated by IZA may include views on policy, but the institute itself takes no institutional policy positions.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit company supported by Deutsche Post World Net. The center is associated with the University of Bonn and offers a stimulating research environment through its research networks, research support, and visitors and doctoral programs. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

Marriage, Wealth, and Unemployment Duration: A Gender Asymmetry Puzzle*

This note presents evidence of the following gender asymmetry: the job-finding effort of married men and women is affected by the income of their spouses in opposite directions. For women, spouse income influences job finding negatively, just as own wealth does: the more the man earns and the wealthier the woman is, the longer it takes for her to find a job. The contrary is the case for men, where spouse income affects job finding positively: the more the wife earns, the faster the husband finds a job. This is so despite the fact that greater own wealth also prolongs unemployment spells for men. These findings are hard to reconcile with the traditional economic model of the family.

JEL Classification: D1, J4, J6

Keywords: gender asymmetries, wealth effects on job finding, unemployment duration

Corresponding author:

Rasmus Lentz
Department of Economics
Boston University
270 Bay State Road, Room 403
Boston, MA 02215
USA
Email: rlentz@bu.edu

* We wish to thank Mette Ejrnæs, Michael Svarer, and Chris Taber for comments and suggestions.

1 Introduction

It is relatively well understood how wealth can influence search effort and reservation wages of unemployed individuals and thereby the duration of unemployment. A risk-averse unemployed worker is expected to be less eager to move back into employment the more wealthy he or she is. In the absence of full unemployment insurance, wealth acts as supplementary self-insurance against unemployment: the more wealth one has, the better insured one is in the sense that one can smooth consumption better between spells of employment and unemployment. Greater wealth means less utility difference between being employed and unemployed. Therefore, if job search is costly, a wealthy unemployed worker is less eager to return to employment, and consequently wealthier jobless workers set higher reservation wages as in Danforth (1979) and/or search less as in Lentz and Tranæs (2005). In both cases the prediction is that wealthier workers will experience longer unemployment spells.

But how does the spouse's income fit into this picture? The most straightforward approach would consider spouse income to be a return on an asset, and hence it should influence one's behavior in the same direction as one's own wealth. It is hard to imagine, at least initially, that spouse income would influence one's behavior in the opposite direction to one's own wealth. Either one's spouse's income does not add to one's wealth, or it adds something positive – large, small, or even very small, but positive. However, it seems that this logic applies only to married women. This may appear strange; nevertheless, it is our conclusion after having taken the question to the data. For married men, the opposite is in fact the case: the higher the woman's income is, the faster the husband finds a job. This is puzzling because both unemployed men and women react to their own wealth as expected; greater wealth implies longer unemployment spells. But when it comes to spouse income, the man seems to be an anomaly.

These findings are difficult to explain within a standard economic model of the family. In the last section of this note we discuss this further, but without being able to find a satisfactory explanation. To get a theoretical grip on labor market phenomena related to married workers, such as the ones we discuss here, there is probably no way around viewing labor market behavior and optimal allocation of resources within the family as one big decision problem. By publishing this

note we wish to encourage more rigorous analysis along these lines, and eventually further empirical scrutiny of the gender asymmetry we present below.

There are other gender asymmetries in labor and family economics, but most we know of are directly or indirectly related to child-bearing: for instance, the asymmetric labor supply response to childbirth, where the father increases his labor supply and the mother reduces it. In that type of case it is not possible to reject the theory that it simply reflects a rational response by the family, as a unit, to the fact that childbirth implies that the woman is less effective than the man in the labor market, for a while at least. Whether this period is long or short depends on the birth itself, length of subsequent period of nursing, etc. The asymmetry we present here is not related to childbirth, and an explanation of the puzzle might have to rely partially on assumptions regarding historical traditions and institutions.

2 Wealth, Spouse Income, and the Duration of Unemployment

This section presents empirical evidence on how unemployment duration is affected by wealth and spouse income. We begin with a discussion of our hypothesis.

2.1 The Hypothesis

We are particularly interested in the behavior of unemployed married individuals. Being unemployed, the individual has no earned income. The possible means of self-insurance in this case are wealth (and thus capital income) and the income of the spouse.

A priori, we would expect own wealth and spouse income to influence job search in the same direction, though possibly with different degrees of intensity. The income of a spouse should play qualitatively the same role as one's own wealth in terms of providing self-insurance against unemployment. Both wealth and a spouse who can earn an income are assets from which one could expect a return, most probably at different 'rates of interest'.¹ In the extreme case where there are no links between the two individuals in the marriage, no economies of scale, no common liabilities, etc., we would expect spouse income to have no behavioral effect whatsoever. The other extreme is where the married couple comprises one decision-making unit with joint ownership of

¹Also, in both cases 're-investments' might be needed to keep this going.

all resources, so that they therefore carry out joint optimization. In this case, additional income earned by either of the two individuals simply means greater common wealth, in which case there is no difference between own and spouse income and wealth. Hence, we would expect the same reaction both qualitatively and quantitatively to greater own wealth and to higher spouse income. In between these two extreme cases are models of varying degrees of common ownership where we would expect the same reaction qualitatively but not quantitatively, as the effect of the one spouse being better off might have only a marginal effect on the situation of the other spouse – but still a positive effect.

The conclusion is that we expect greater own wealth and higher spouse income to influence the choice of search effort and reservation wage in the same direction, and therefore also unemployment duration in the same direction. Given the theory mentioned in the introduction we expect wealth and spouse income to have a negative effect on unemployment duration, the reason being that we expect better (self-)insurance to make unemployed workers more patient in waiting for a better job and/or search less intensively in an environment where job search is costly. In the next section we present the data we use to test this hypothesis.

2.2 Data

The analysis is based on Danish micro data from the period 1980-1994.² The data consist of 12,865 unemployment spells ranging from 4 weeks to 52 weeks. The unemployment spells are observed on a weekly basis. Spells of longer duration were censored at 52 weeks. Spells of less than 4 weeks were dropped from the data set. These brief spells do not represent unemployed search but rather short vacations and exogenously-given mismatches between the end of one job and the beginning of another where the new job was found via on-the-job search. Including the less-than-4-week spells in the analysis did not change the point estimates significantly but lowered the statistical significance of the estimates somewhat. The data only cover workers aged between 18 and 66.

The unemployment spell data are then merged with administrative databases that provide demographic and financial information about the workers associated with the unemployment spells.

²The data set was generously made available by The Centre for Labour Market and Social Research in Århus, Denmark.

Table 1: Summary Statistics

Variable	Minimum	Maximum	Mean	Standard Deviation
Years of education	9.00	18.00	11.99	2.64
1=Female	0.00	1.00	0.47	0.50
Spouse's income	-159,760.00	886,239.00	0.00	137,138.00
# of children	0.00	12.00	1.66	1.88
Age	18.00	66.00	42.51	8.86
1=Upper management	0.00	1.00	0.15	0.35
1=Lower management	0.00	1.00	0.21	0.40
1=Salaried worker	0.00	1.00	0.29	0.45
1=Skilled worker	0.00	1.00	0.12	0.32
Wage	56.78	925.53	146.24	59.74
1=Owner of real estate	0.00	1.00	0.53	0.50
Net wealth	-1,046,000.00	1,046,000.00	115,313.00	307,580.00

N=12,865

All income and wealth amounts are in 1994 DKK.

This information is observed on a yearly basis. Thus, changes in a worker's wealth and spouse income are not observed during an unemployment spell. The variables capture the resources the unemployed worker has for self-insurance at the beginning of his or her unemployment spell. More specifically, the wealth and income of an individual are measured on the last day of the year prior to the year of the unemployment spell, and spouse income is the income during the year of the unemployment spell. The data are summarized in table 1. The residual category related to the occupation dummies is unskilled workers. The spouse's income is defined as the difference from the mean spousal income. In the next section we present the model we use to test this hypothesis.

2.3 The Empirical Model

We wish to determine the impact of wealth holdings and spouse's income as well as a range of other co-variates on the hazard rate of unemployment for individual workers. The analysis assumes a basic proportional hazard model similar to Meyer (1990). Thus, the individual hazard rate is assumed to take the following form,

$$\lambda_i(t) = \lambda_0(t) \exp(z_i' \beta), \quad (1)$$

where $\lambda_0(t)$ is the baseline hazard common to all individuals. The expression (1) is assumed to be continuous over time. Note also that the individual co-variates z_i are constant over the unemploy-

ment spell; this is dictated by the data. Since observations of workers leaving unemployment are made only on a weekly basis rather than continuously, it is not known at exactly which point in time a worker leaves the unemployment pool. Let t_i denote the duration of worker i 's unemployment spell. Meyer (1990) derives the log-likelihood expression for this case, which is given by:

$$\log L(\gamma, \beta) = \sum_{i=1}^N \left[\delta_i \log \{1 - \exp[-\exp(\gamma(t_i) + z_i' \beta)]\} - \sum_{t=1}^{t_i-1} \exp[\gamma(t) + z_i' \beta] \right], \quad (2)$$

where $\delta_i = 0$ denotes that spell i is right censored and $\delta_i = 1$ that the spell ending was observed. Furthermore, $\gamma(t)$ is defined by:

$$\gamma(t) \equiv \log \left(\int_t^{t+1} \lambda_0(s) ds \right), \quad t = 1, \dots, T-1. \quad (3)$$

One can interpret this model as a discrete model, in which case $\gamma(t)$ is the estimated baseline hazard at time t . Alternatively, the model can be interpreted as a continuous time hazard model, where the data do not allow for identification of the continuous baseline hazard. There will be an infinite number of baseline hazard functions $\lambda_0(\cdot)$ which will satisfy (3).

Given the fact that the effect of a given worker characteristic on the hazard rate is identified mostly via cross-sectional data, one might be worried that unobserved heterogeneity could bias the results. To control for this one can add an unobserved heterogeneity term to the hazard rate. Furthermore, in this procedure one can exploit the fact that most workers experience multiple spells of unemployment. Let t_{ij} denote the duration of the j 'th spell for worker i . Similarly let z_{ij} be the set of co-variates for the j 'th spell of worker i . Individual i 's unemployment hazard for the j 'th spell is then given by:

$$\lambda_{ij}(t) = \exp(\gamma(t) + z_{ij}' \beta + \theta_i),$$

where θ_i is the unobserved heterogeneity term. The analysis will follow the Heckman and Singer (1984) approach and assume that θ_i follows a discrete distribution which will be non-parametrically estimated. Let L be the number of support points of the distribution and let J_i be the total number of unemployment spells of worker i . Thus, the likelihood of observing all of worker i 's spells is given

by:

$$L_i(\gamma, \beta, \theta, \Pr(\theta)) = \sum_{l=1}^L \Pr(\theta_l) \exp \left\{ \sum_{j=1}^{J_i} \left[\delta_{ij} \log \{1 - \exp[-\exp(\gamma(t_{ij}) + z'_{ij}\beta + \theta_l)]\} - \sum_{t=1}^{t_{ij}-1} \exp[\gamma(t) + z'_{ij}\beta + \theta_l] \right] \right\}. \quad (4)$$

The full log-likelihood for the data set is then given by:

$$\log L(\gamma, \beta, \theta, \Pr(\theta)) = \sum_{i=1}^N \log L_i(\gamma, \beta, \theta, \Pr(\theta)). \quad (5)$$

The distribution of the unobserved heterogeneity term will be estimated following the approach set out in Heckman and Singer (1984). The number of support points in the distribution is increased until the estimation is no longer improved by additional points.

2.4 Results

Results based on both (2) and (5) are reported in table 2. The set of co-variates includes all of the characteristics summarized in table 1. Also included is a set of yearly dummies to account for any year-to-year differences in the hazard rate. We allow for gender differences by the inclusion of multiplicative terms, the main one being for spousal income and an additional one being for the number of children. The baseline hazard estimate and the yearly dummy estimates for both specifications are displayed in figures 1 and 2.

First of all, it is seen that the wealth effect is negative, as expected, but also quite small. In the analysis with no unobserved heterogeneity it is highly statistically significant. Once unobserved heterogeneity is included, the estimate remains negative and is of the same magnitude, but the statistical significance of the sign of the effect is now only at the 10% level.³ This still suggests a fairly significant effect. The absolute effect is quite small which is consistent with the fact that most duration studies find either negative or no duration dependence of the baseline hazard. If the wealth effect were large one would expect to see signs of positive duration dependence. Similar results have been established using Dutch and French data in Bloemen and Stancanelli (2001) and Algan, Chéron, Hairault, and Langot (2001) respectively.

³A test of the hypothesis that the wealth effect is negative against the alternative that it is not.

Table 2: Results of Proportional Hazard Estimation

	Without Unobserved Heterogeneity		With Unobserved Heterogeneity	
	Point Estimate	Std. Dev.	Point Estimate	Std. Dev.
Years of education	-0.0015	0.0085	-0.0415	0.0577
1=Female	0.0118	0.0211	0.0147	0.0443
Age	-0.4336**	0.1194	-0.7045**	0.1742
Spousal income	2.89e-7**	1.30e-7	4.08e-7**	1.66e-7
(1=Female)*(Spousal income)	-4.62e-7**	1.59e-7	-6.48e-7**	2.04e-7
# of children	0.0490	0.0712	0.0568	0.0998
(1=Female)*(# of children)	-0.1118	0.0993	-0.0116	0.1454
1=Upper management	0.0111	0.0385	-0.0005	0.0410
1=Lower management	0.0358*	0.0293	0.0584*	0.0374
1=Salaried worker	0.0279	0.0257	0.0150	0.0367
1=Skilled worker	0.0298	0.0359	0.0103	0.0453
Wage	4.91e-4	4.27e-4	6.47e-4	5.62e-4
Wage ²	-8.96e-7*	7.07e-7	-1.06e-6	9.23e-7
1=Owner of real estate	-0.0289	0.0211	-0.0128	0.0252
Net wealth	-8.86e-8**	3.41e-8	-5.49e-8*	4.49e-8
θ_1			-0.5617	0.0511
θ_2			0.0000	-
θ_3			0.7333	0.0632
θ_4			1.4639	0.1360
Pr(θ_1)			0.2665	0.0565
Pr(θ_2)			0.4532	0.0642
Pr(θ_3)			0.2411	0.0339
Pr(θ_4)			0.0393	0.0123

*The sign of the estimate is significant at the 10% level.

**The sign of the estimate is significant at the 5% level.

There is a highly significant difference between men and women with respect to the effect of spousal income: the higher his spouse's income, the shorter a man's unemployment spell will be. This is in contrast to the effect of wealth. For women, the effect is the direct opposite. The coefficients for women alone are $-1.73e-7$ and $-2.40e-7$ without and with unobserved heterogeneity respectively, and the corresponding standard deviations are $0.67e-7$ and $1.26e-7$. Thus, the effect is strengthened when we control for unobserved heterogeneity, where the result is significant at the 6% level (figures not reported in the table).

As discussed above, one would expect a negative coefficient to spouse income – as we get for women – because the insurance associated with the spouse's income should make the unemployed worker search with less intensity and/or want to set a higher reservation wage.

The number children is not significant when we control for unobserved heterogeneity for either gender.

Other notable results include the negative effect on the hazard rate of age. As workers get older, the unemployment hazard rate decreases. The effect of education is also negative but not significant. Both of these signs are consistent with the results in Meyer (1990). Generally, once one controls for the wage level, then theory gives little guidance as to the signs of the effects of education and occupation. And indeed only in the case of lower management do the estimates show some significance.

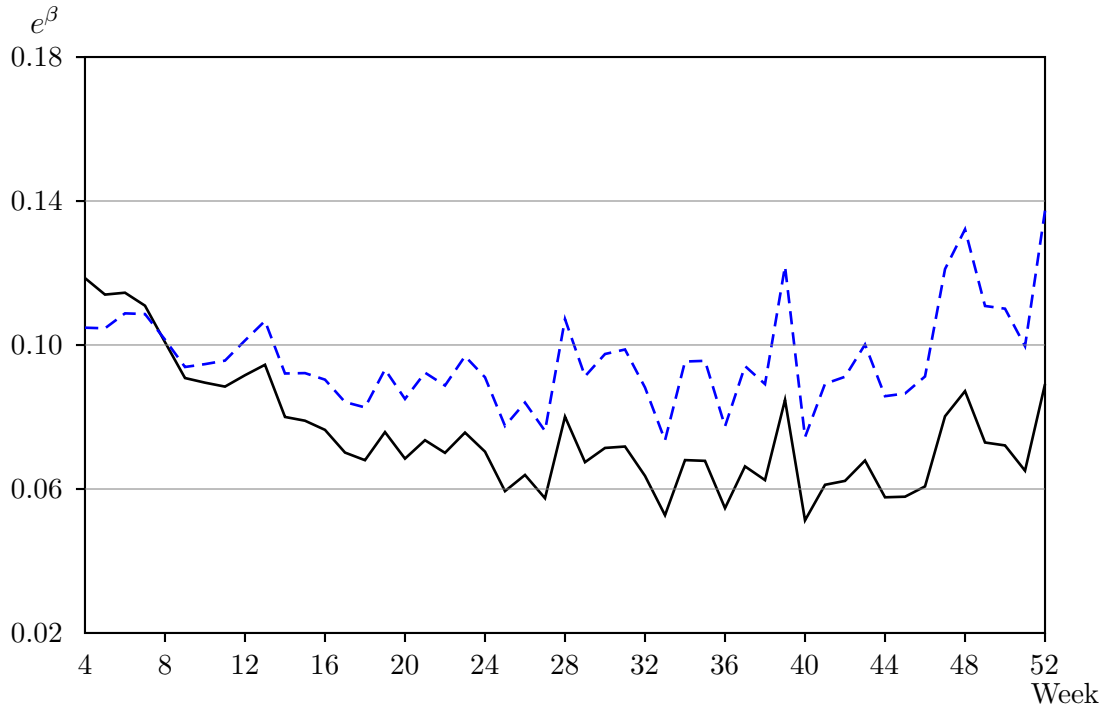
The dummy variable for whether or not the worker owns real estate is included in the analysis in order to pick up on any possible liquidity effects. Essentially, one might be worried that the worker's wealth might not have the liquidity required to work as a vehicle for consumption smoothing if a sizeable proportion of it were tied up in a house. Of course, if capital markets are perfect, it should be possible to borrow against the value of real estate, and consequently the degree of liquidity of assets should not matter. However, it is not clear that the Danish capital markets were sufficiently complete during the period studied to remove any such liquidity effects. One would expect the sign of the estimate to be positive if workers could not borrow against the value of their real estate. The actual estimate is negative but not significantly so. Thus, one cannot draw any strong conclusions as to whether liquidity matters or not.

The estimation allows for a non-linear wage effect. Lentz (2002) shows that a job search model with savings can produce a bell-shaped relationship between the unemployment hazard rate and the wage.⁴ Our estimation does in fact produce such a relationship and it is robust to the inclusion of unobserved heterogeneity. At lower wage levels, an increase in the wage will increase the unemployment hazard rate but for higher wage levels the relationship is reversed.

The baseline estimates are shown in figure 1. It shows that the estimates with no unobserved heterogeneity display some negative duration dependence for the first 26 weeks. For greater durations the baseline estimate is more or less constant. When unobserved heterogeneity is included the baseline is more or less constant, with some suggestions of a U-shape. Clearly, the slight upward

⁴See Lentz (2002) for a detailed description of this relationship.

Figure 1: Baseline Estimate.



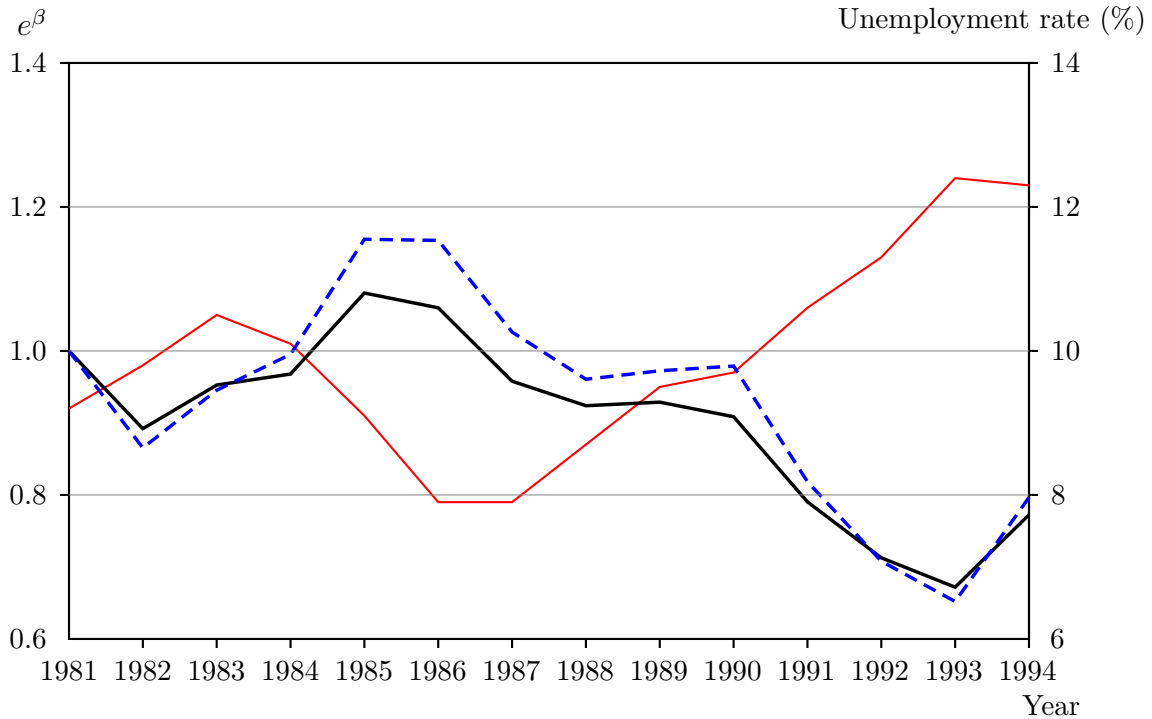
Note: The point estimate for the estimation without unobserved heterogeneity is drawn in solid pen and the estimate with unobserved heterogeneity is drawn in dashed pen.

swing towards the end is due to the selection effect associated with the unobserved heterogeneity.

Finally, the yearly dummies closely follow the aggregate unemployment rate with a one-year lead. Abowd, Corbel, and Kramarz (1999) find that firms mainly adjust their hiring behavior to fit their employment needs and maintain a fairly constant separations rate. Thus, in a recession the offer arrival rate goes down and consequently the duration of periods of unemployment increases. At a constant or higher flow into unemployment the unemployment rate will increase, producing the relationship seen in the estimation.

We conclude that while the two genders react symmetrically with respect to wealth in that wealthier workers experience longer unemployment spells, men and women react asymmetrically with respect to spouse income. It looks as if spouse income is like own wealth to the married women when it comes to insuring against unemployment; to the married man, however, spouse income complements their own wealth and thus reduces – not increases – unemployment duration. Thus, women seem to behave as expected: it is the men that are the anomaly.

Figure 2: Yearly Dummy Estimates.



Note: Yearly dummy estimates on left axis. The point estimate for the estimation without unobserved heterogeneity is drawn in solid bold and the estimate with unobserved heterogeneity in dashed bold. The Danish unemployment rate is drawn on the right axis in thin pen.

3 Discussion

When we have presented the spouse-income gender asymmetry reported above to social-science scholars, we have been offered two types of explanation for the phenomenon. As mentioned in the introduction, we are not convinced by any of these explanations, and we hope that this note will inspire further work on the question.

One type of explanation starts by assuming that the two genders have different preferences and then tries to explain the formation of these differences historically, and possibly also as part of a bigger explanation that encompasses the formation of other institutions in society. Once asymmetric preferences are accounted for, the phenomenon discussed here follows suit. The other type of explanation begins with symmetric preferences and then derives an explanation from differences in endowments and the division of labor at the labor market and within the family. This latter type

of explanation is in line with the traditional economic approach.

To the first type belongs an explanation based on the claim that men see themselves as the breadwinners and the masters of the house, which for some reason is not the way women see themselves. Presumably the man then finds his position threatened if he does not work while his wife is providing, and the more she provides the more threatened he feels. Explanations along these lines are often suggested to us when we present the spouse-income gender asymmetry at seminars. There does exist empirical evidence that is in accordance with such "breadwinner syndrome" or "master-of-the-house syndrome" hypotheses. E.g., Weiss and Willis (1997) finds that marriages are more likely to be dissolved after the woman experiences an unanticipated earnings increase and less likely to dissolve when the man's income surprise positively (see also Svarer (2002)). Of course, this does not bring us any closer to an explanation of the gender asymmetry we discuss here. The assumption of preference heterogeneity obviously facilitates the explanation of heterogeneous behavior in economics, but the economic literature does not offer any clues as to why the reactions should depend systematically on gender.

An explanation of the second type would begin by assuming that the man and woman are endowed very differently and that the value to the family of one or the other of them being home is very different as a consequence. The problem here is that the value to the family of the man being at home would have to be strictly negative before we could expect the signs we find empirically, as we would expect the wealth effect of higher spouse income to be present as well, but overpowered by this (negative) productivity effect. And this is a strong assumption, not much different really from assuming that men and women simply have different tastes and aspirations, as assumed in the first type of explanation.

A third type of explanation, using elements from both the preference and the endowment approaches, is one that assumes that the incomes of the two spouses do not enter into the family wealth symmetrically, either formally by contract (as happens in certain developing countries) or implicitly, in which case we approach the first type of explanation. So the income of the man is family income, whereas the income of the women is her own. Again, this is an assumption that more or less implies the result directly.

Our interpretation of this is that one has to model labor market behavior and resource allocation within the family as one joint decision problem before one can hope to produce satisfactory explanations. Unfortunately, this avenue is underdeveloped in the theoretical literature.

References

- Abowd, John M., Patrick Corbel, and Francis Kramarz (1999). The entry and exit of workers and the growth of employment: An analysis of french establishments. *Review of Economics and Statistics* 81, no. 2: 170–87.
- Algan, Yann, Arnaud Chéron, Jean-Olivier Hairault, and Francois Langot (2001). Wealth effect on labor market transitions. *Working Paper*.
- Bloemen, Hans G. and Elena G. F. Stancanelli (2001). Individual wealth, reservation wages, and transitions into employment. *Journal of Labor Economics* 19, no. 2: 400–39.
- Danforth, John P. (1979). On the role of consumption and decreasing absolute risk aversion in the theory of job search. In S. A. Lippman and J. J. McCall (Eds.), *Studies in the Economics of Search*, pp. 109–31. New York: North-Holland.
- Heckman, James J. and Burton Singer (1984). A method for minimizing the impact of distributional assumptions in econometric models for duration data. *Econometrica* 52, no. 2: 271–320.
- Lentz, Rasmus (2002). Optimal unemployment insurance in an estimated job search model with savings. *Working Paper*.
- Lentz, Rasmus and Torben Tranæs (2005). Job search and savings: Wealth effects and duration dependence. *forthcoming in Journal of Labor Economics*.
- Meyer, Bruce D. (1990). Unemployment insurance and unemployment spells. *Econometrica* 58, no. 4: 757–82.
- Svarer, Michael (2002). Determinants of divorce in denmark. *University of Aarhus Working Paper 2002-19*.
- Weiss, Yoram and Robert J. Willis (1997). Match quality, new information, and marital dissolution. *Journal of Labor Economics* 15, no. 1: S293–329.