Institut für Höhere Studien
Institute for Advanced Studies
Vienna

Continuing Education Center

MSc Economics

Unemployment and Non-Participation: History Dependence in Job Finding Probabilities

A Master's Thesis submitted for the degree of
"Master of Science"

supervised by
Tamás K. Papp

Bence A. Bardóczy

1325887

Vienna, 2015

MSc Economics

Affidavit

I, Bence András Bardóczy

hereby declare

that I am the sole author of the present Master's Thesis,

Unemployment and Non-Participation: History Dependence in Job Finding Probabilities

46 pages, bound, and that I have not used any source or tool other than those referenced or any other illicit aid or tool, and that I have not prior to this date submitted this Master's Thesis as an examination paper in any form in Austria or abroad.

Vienna, June 8, 2015

Signature

Acknowledgements

I would like to thank Tamás Papp for guiding me throughout the thesis and Christian Haefke for providing me with longitudinally matched CPS data. Their help was essential for the successful completion of this project. Furthermore, I am grateful to Fabian Greimel, Tímea Virágh and Leopold Zeßner-Spitzenberg for valuable comments and discussions.

Contents

List of Figures

List of Tables

Abstract

I use a novel empirical approach to study heterogeneities in labor market flows in the U.S. Using data on short-term labor force history, I develop Bayesian logit models that capture large variations in the job finding probabilities not only of the unemployed but of the non-participants, as well. A decomposition of these variations by time periods suggest that recent employment has a bigger impact on current job finding probabilities than contemporaneous search behavior. Furthermore, I document that among prime age men, non-participants are almost as likely to start working as active job-seekers. The gap narrows even further in recessions, for the job finding probability of the unemployed-looking tends to fall disproportionately.

1 Introduction

Losing one's job or being unable to find one are arguably the biggest idiosyncratic risks that most individuals have to face, especially in recessions. Consequently, the unemployment rate is one of the most widely followed economic statistics, featured in the popular press just as often as in academic journals. Importantly, movements in labor market stocks emerge from the continuous transitions between labor force states. For example, unemployment might rise in recessions because many more workers get fired or because jobs are hard to find. Therefore, describing the interplay and cyclical properties of worker flows is crucial for understanding labor market fluctuations.

The majority of today's structural models of labor market flows belongs to the Diamond-Mortensen-Pissarides (DMP) framework as described for example by Pissarides (2000). These models focus on the frictional job search of unemployed workers and usually abstract from non-participation. However, there are at least three good reasons for introducing non-participation into flows-based models. First, the definition of labor force participation itself is ambiguous. It has long been debated for example, whether discouraged workers—who want to work but are not searching because they believe that they could not find a job—should be counted unemployed or not.[1] Second, as Figure 1 shows, more non-participants have started to work in every month in the last 25 years than unemployed. I have the impression that this fact is not widely recognized, even though it is a powerful argument for studying non-participants. The third line of reasoning, put forward recently by Elsby, Hobijn and Şahin (2013), highlights the importance of cyclical movements between unemployment and non-participation in shaping the unemployment rate.

As a consequence, there is a growing interest in developing a unified model of aggregate employment, unemployment and participation. Practically, this means the inclusion of an operative labor supply channel to modern flows-based models with search frictions. Significant advances have been made in this direction by Krusell, Mukoyama, Rogerson and Şahin (2010, 2011, 2012). Non-participants however, constitute a large and widely heterogeneous fraction of the population. A prime age man who suspends active job-search for a while is classified as being out of the labor force just as retired or permanently disabled persons are. Any attempt to formalize the behavior of non-participants should deal with this heterogeneity. A difficult task, as our knowledge of the key aspects of non-participation is still rather limited. This lack of tangible stylized facts calls for further empirical studies on flows in- and out of the labor force.

Our understanding of unemployment is much more advanced, and has the potential to inform our approach to non-participation. The key aspect of unemployment is active job search. Search requires effort, hence needs to be incentivized. The incentive is the prospect of earning a surplus once successfully matched with an employer. The size of the surplus, hence optimal search effort and the implied job finding probability, depends on macroeconomic conditions and individual traits, as well. In most DMP models, the

[1] See Jones and Riddell (1999) and references therein.

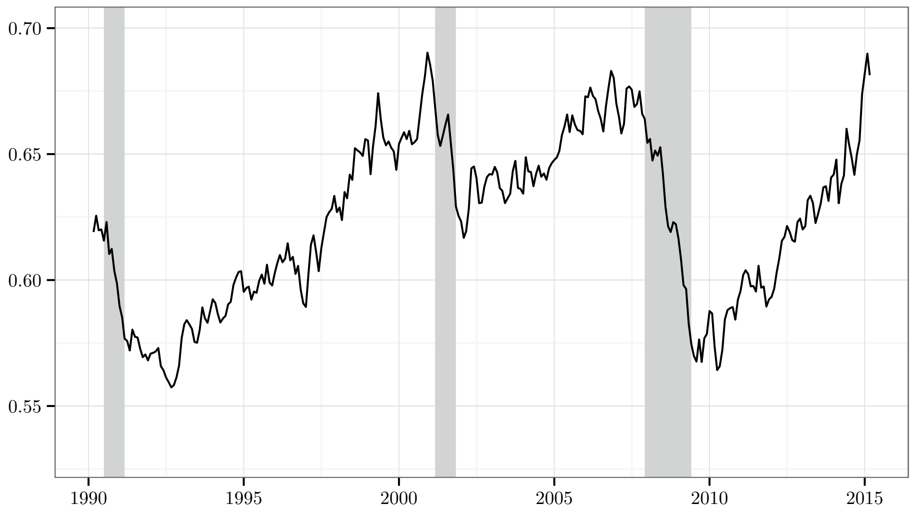

Figure 1: Share of Non-Participants among New Inflows to Employment

Note: Three-month moving averages of seasonally adjusted monthly flows. Re-weighted to be representative of the US population. Official BLS series, downloaded from FRED.

relevant aggregate state variable is labor market tightness, the ratio of vacancies and unemployed in the economy. Unemployment and vacancies are negatively correlated and volatile in the data, thus job finding probability is cyclical due to macro fluctuations that affect all job-seekers.

At the individual level, job finding probability is observed to fall with unemployment duration. Regardless of its reasons, such heterogeneity contributes to the procyclicality of average job finding probability as illustrated by Figure 2. Subfigure 2a shows that job finding probabilities fell almost uniformly for all duration from 2006 to 2009 (i.e. from peak to trough of the Great Recession). Apparently, negative duration dependence in the job finding probability of the unemployed is a salient feature of the labor market. Turning to subfigure 2b reveals that the common decline was accompanied by an unequivocal shift towards longer unemployment duration. Thus, one can infer that the unemployment rate would have had risen to some extent due to this compositional shift alone.

Exploration of how aggregate- and compositional effects jointly shape labor market flows throughout the business cycle is the bigger research agenda underlying this master's thesis. In the next subsection, I briefly review the academic papers that directly influenced my work.

1.1 Literature Review

It has long been observed that the unemployed have a harder time finding jobs in recessions. A classical explanation for the procyclicality of job finding probability, put forward by Darby, Haltiwanger and Plant (1985, 1986), points to the role of worker heterogeneity. They conjectured that there are two types of unemployed workers. First, a high-turnover group who undergoes frequent but short unemployment spells, i.e.

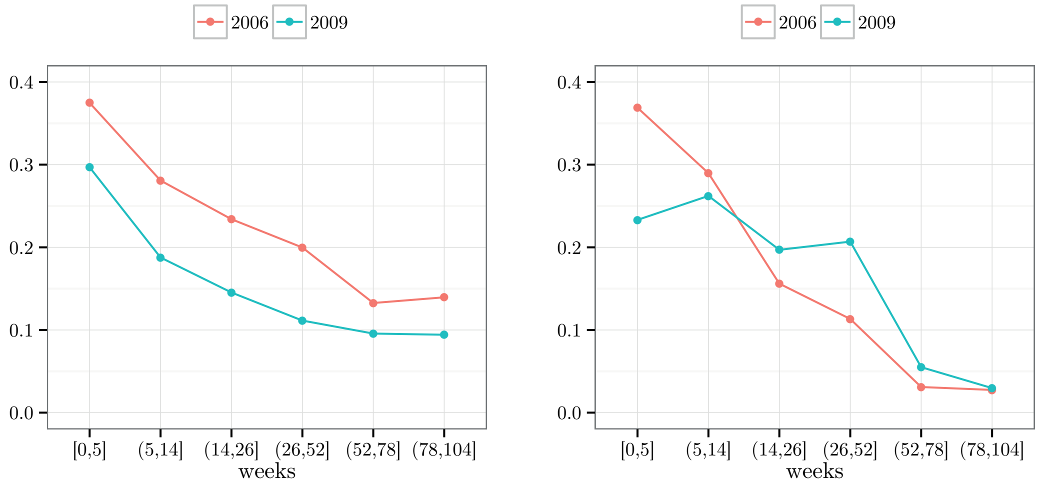

(a) job finding probability

(b) share in sample

Figure 2: Unemployment Duration and Job Finding

Note: Unadjusted CPS micro data for unemployed aged 16-65. Includes records matched across two months from rotation groups 2-4 and 6-8.

has high job finding probability. Second, a low-turnover group whose members seldom become unemployed, but then need more time to find a suitable job, thus their measured job finding probability is low. In recessions, disproportionately many "permanent jobs" get terminated, thus the share of low-turnover unemployed increases. All in all, unemployment rises primarily due to job destruction, while there is a side-effect: a compositional shift that lowers the average job finding probability and aggravates unemployment.

In a very influential paper, Shimer (2012) argues against this "heterogeneity hypothesis". He develops a continuous time unemployment accounting framework with two states—employment, unemployment—to measure job finding and separation rates. His method requires only aggregate data on employment and (short-term) unemployment, which makes it robust compared to standard measures of worker flows that rely on survey data which is noisier. His main point is that the job finding rate displays stronger cyclical patterns than the separation rate, hence contributes more to unemployment fluctuations. As a robustness check, he introduces non-participation as a third state, which increases the number of transition rates to six. The U–E transition rate retains its prominent role, but the elegant accounting model does not apply anymore. Thus, the first strong assumption he makes is abstracting from non-participation, as it was commonly done in the macro-labor literature until recently.

The second strong assumption behind Shimer (2012)'s method is the homogeneity of transition rates. In any given month, all the unemployed are assumed to have the same probability of finding a job, which is determined solely by aggregate shocks. He supports this assumption by means of the following exercise. First, he divides the unemployed into different groups.[2] He shows, that in the presence of such heterogeneity, his measure of the job finding rate gives the average job finding rate for all the unemployed. Second, he decomposes the variance of the average job finding rate to changes in population shares of the groups and changes in the group specific job finding rates. He finds that composition effects explain only a minor part of the variance. He concludes that homogeneity is an acceptable approximation, thus Darby, Haltiwanger, and Plant were wrong on two accounts: falling job finding is a more important source of high unemployment than spikes in job destruction, and worker heterogeneity plays a negligible role compared to changes in the macro conditions. The caveat in this approach is that workers might differ in ways which the covariates do no identify. Consequently, sizable composition effects with respect to unobserved, or simply omitted, heterogeneity cannot be ruled out.

Omitted unemployment duration is central to the criticism of Hornstein (2012), who warns that the duration distribution of unemployment implied by homogeneous transition probabilities is inconsistent with the data. He refines Shimer's 2-state model to distinguish between short- and long-term unemployment. He finds that the transition rates of the long-term unemployed are more volatile than that of the newly unemployed. Although job finding rate remains dominant, the employment exit rate of the long-term unemployed too has a strong correlation to the unemployment rate. His estimations suggest that the observed negative duration dependence in job finding is not causal. Rather, reported unemployment duration is more likely to pick up unobserved heterogeneity. This means that the long-term unemployed have low job finding probabilities, because they are less employable in the first place. Conversely, people who find jobs quicker would be more likely to find jobs after long spells of unemployment, too. Easily employable people do not have to search for long, thus only people with inherently low job finding probability get stuck in long-term unemployment.

Elsby, Hobijn and Şahin (2013) also stress the importance of worker heterogeneity. They use a 3-state model—employment (E), unemployment (U), and non-participation (N)—and find that flows on the participation margin are notably cyclical and account for one-third of unemployment fluctuations. They argue that part of the story is explained by compositional effects. In recessions, more workers with strong labor market attachment become unemployed. These workers are less likely to abandon search and quit the labor force, hence the average U–N transition probability falls, which increases the unemployment rate. What is interesting about their results is that they revive the heterogeneity hypothesis of Darby, Haltiwanger and Plant by a methodology inspired by Shimer. The explanation of this difference is twofold. First, as Hornstein pointed out, Shimer did not consider duration dependence. Thus, he missed an important ingredient of worker heterogeneity. In addition to the variables used by him, Elsby, Hobijn and Şahin include labor force status a year before. Moreover, they interacted all the variables, thus accounted for compositional shifts between substantially more

[2] He considers gender, age, marital status, education, census division, reason for unemployment. One at a time, without interactions.

groups than Shimer, who isolated heterogeneity in education, marital status etc. Second, if the unemployed with lower U–N transition rates also had higher job finding rates, then the 2-state accounting scheme would understate the compositional change in unemployment exit rates: higher U–E and lower U–N transitions offset when are viewed together as the "outs of unemployment".

Shimer (2012); Hornstein (2012); Elsby et al. (2013) all agree that the job finding probability is procyclical and is strongly correlated with the unemployment rate. However, they base their analysis on variance decomposition, not structural models. Thus, there is no guarantee that they identify causal relations. Coles and Moghaddasi (2014) consider a standard DMP model with homogeneous workers, where vacancy creation is inelastic and productivity and job destruction shocks are negatively correlated.[3] When a job destruction shock raises unemployment, jobless workers essentially start to deplete the vacancy stock and the job finding probability falls persistently. Although the correlation of unemployment and job finding probability is large, it is triggered by job destruction. The model's predictions match the main findings of Shimer (2012) yet support the insight of Darby, Haltiwanger, and Plant that unemployment volatility is mainly driven by job destruction shocks.

Kroft, Lange, Notowidigdo and Katz (2014) impose a stylized search and matching structure on a 3-state model with negative duration dependence in the job finding probability of the unemployed. Although theirs is not a structural model either, it is certainly a step closer than variance decomposition studies. The authors find that allowing for duration dependence in the job finding probability and cyclical movements on the participation margin are both needed to match the unprecedented rise in long-term unemployment and the outward shift of the Beveridge curve observed in the Great Recession. However, their model gives rather poor predictions of the participation rate and the job finding probability of non-participants. I conjecture that these shortcomings originate from their treating the non-employed asymmetrically. Whereas they model negative duration dependence in the job finding probability of the unemployed, they assume that non-participants find jobs at a homogeneous rate. The only reason I can think of for making this assumption is data limitation: in the Current Population Survey (CPS) only unemployed respondents are asked about the length of current unemployment spell, similar data is not collected for non-participants.

The papers reviewed so far suggest that participation decisions and negative duration dependence in job finding are important ingredients of cyclical phenomena on the labor market. Another line of research addresses the controversial nature of drawing a line between unemployment and non-participation. The purpose of the classification system is to capture labor market attachment of the non-employed. Jones and Riddell (1999, 2006) argue that instead of relying on self-reported job-search behavior, clas-

[3]The usual free entry assumption makes vacancy creation perfectly elastic. With more unemployed searching for jobs, it is easier for firms to fill vacancies, hence ceteris paribus vacancies become more valuable and firms will create more of them. As a result, unemployment and vacancies are positively correlated in the standard DMP model, a strongly counterfactual prediction. See Shimer (2005), Hornstein, Krusell and Violante (2005) and Coles and Moghaddasi (2014).

sification should be "evidence-based". What they mean by this is that labor market attachment of a group is revealed by its transition probabilities to other states. They start out from a finer partition of the non-employed, and estimate logit models of individual transitions using micro data from the Canadian Labor Force Survey. Controlling for largely the same variables as Shimer (2012), they find that those who want a job, but are not actively searching, have significantly different transition rates both from the unemployed and from other non-participants. They also note that workers on temporary layoff have much higher job finding probability than the unemployed-looking. Importantly, the authors call attention to the fact that some non-participants display rather high labor market attachment.

1.2 Goals for the master's thesis

Building on the insights outlined above, I undertake a novel empirical analysis of heterogeneities in labor market states. Elsby et al. (2013) argued that prior labor force status reveals differences in labor market attachment that could not be captured by contemporaneous observables. They find sizable composition effects in the U–N transition rate by distinguishing the currently unemployed based on their labor force status (E, U, or N) a year before. I elaborate on this idea by using data on short-term labor force history and apply it to job finding probabilities. Specifically, I consider three-month histories preceding potential transitions to employment. Job finding probability has large fluctuations at business cycle frequency, but that does not necessarily imply that it is driven by aggregate shocks. The unemployed constitute a relatively small fraction of the working age population, with relatively large turnover due to flows to employment and non-participation. This means that the composition can change rather rapidly, too.

I analyze history dependence in the context of binomial logit models of job finding probability using longitudinally matched micro data from the CPS. This regression framework is similar to Jones and Riddell (1999, 2006), who estimated both multinomial and binomial logit models of transitions between labor market states. The crucial difference is that they explicitly assume that all but the last status are irrelevant for predicting transition rates. This Markov assumption is implicit in the majority of labor-macro studies, but omitting history from an econometric estimation of transitions might bias the results. In fact, I find that a 3-state model (E, U, N) with three-month history captures more variation in job finding probabilities than a finer 5-state model without history, and that current status and history are correlated.

Intuitively, labor force histories are closely related to duration of current spell. Indeed, I demonstrate that allowing heterogeneity in job finding probability with respect to labor force history is sufficient to match the negative duration dependence stressed by Hornstein (2012) and Kroft et al. (2014). I also claim that looking at history has conceptual as well as practical advantages over duration. Conceptually, a sequence of states contains more information than just the length of the current spell. Of course, it is a higher dimensional object, thus comes at the cost of larger data requirement, but my hypothesis is that this extra information is valuable. Practically, most employment surveys do not gather information on time spent out of the labor force. Therefore, working with labor force history enables researchers to consider "duration dependence" for non-participants as well as for the unemployed.

Treating three-month sequences of labor market status as potentially distinct states leads to the fragmentation of the sample. Some histories, for example NEU, are rather rare thus have few observations. Having to work with small sample sizes motivates my using of Bayesian inference. Conditional on model specification, Bayesian models give precise probability statements for arbitrarily small samples and do not rely on asymptotic arguments. Bayesian methods are also well-suited to fit multilevel models. This flexible framework allows me to quantify the relative importance of status at months $t-1$, $t-2$, $t-3$ for explaining month t transition rates. I find that recent employment has bigger impact than contemporary search-behavior.

The thesis is structured as follows. In Section 2, I describe where the data is coming from and how the sample is constructed. In Section 3, I estimate the baseline models with and without history and compare their predictions. In section 4, I develop a more elaborate multilevel model, and use it to decompose the variation in job finding probabilities by labor force status in the preceding months. Section 5 concludes.

2 Data

I use micro data from the Current Population Survey (CPS), the source of official employment statistics in the U.S. since 1948. In the past 67 years, the CPS has undergone several major revisions, the last of which was made in 1994. Since 2000, approximately 70,000 households are interviewed monthly based on their addresses. The survey is a rotating panel: respondents enter the sample for four months, rotate out for eight months, and then rotate back again for another four months. Therefore in principle, three-quarters of respondents are surveyed in two consecutive months (rotation groups 2-4 and 6-8). The structure of the survey therefore enables direct measurement of worker flows. For example, the average U–E transition probability is the ratio of workers who report being unemployed in one month and employed in the next. Analogously, it is in principle possible to match one-quarter of the records for four consecutive months. Thus, for rotation groups 4 and 8, the complete set of variables is available for three months before an observed potential transition. Of course, there is further loss of data due to survey non-response.

These measures of U.S. labor market flows have been used extensively in academic research. Recent examples include Shimer (2005, 2012); Elsby et al. (2013); Kroft et al. (2014) just among the papers I have already cited. These authors computed (or were provided with) transition probabilities starting from 1967. It is worth noting however, that the Bureau of Labor Statistics (BLS) itself, the institution that administers the CPS, suspended publishing gross flows data in 1952 due to concerns about its accuracy. Eventually, publication of the official series has been resumed, but only for 1990 onwards (see Frazis et al. (2005) and Ilg (2005). In the master's thesis, I only use post-1996 data, which were matched and made available to me by Christian Haefke, whom I thank again. Appendix A contains the official definitions of labor market states that apply for this period, and the notation I use throughout the master's thesis.

2.1 Margin- and classification errors

Before moving on to the analysis, I shortly discuss two types of errors in the CPS that are known to affect measured flows. The first is *margin error*, the inconsistency of measured stocks and flows. While labor market stocks are constructed using all eight rotation groups, flows are necessarily based on the subsample that can be matched across a minimum of two consecutive months. There are three potential sources of inconsistency. First, systematic rotation group effects. Second, non-random attrition due to changes in residency or other reasons. Krueger, Mas and Niu (2014) find that the two are connected: significantly fewer respondents claim to be unemployed in later rotation groups, and the pattern of the bias over time mirrors that of survey non-response. Third, flows in and out of the scope of the CPS (turning 16, emigration, death and so on). One possible way to deal with margin error is to solve for stock-consistent flows with minimal adjustment (in sense of weighted sum of squares). Practically, margin error adjustments have small effect on measured transition probabilities as shown for example by Elsby et al. (2013).

Classification errors, assigning a worker to the wrong labor market state, are potentially more important. While some of these mistakes might offset each other in stock measures, they are likely to generate spurious transitions. For example, if hundred workers with true labor force history UUU are misclassified as UEU and similarly hundred EEE are misclassified as EUE, than all errors offset for stocks but there are two hundred spurious E–U and two hundred spurious U–E transitions. Abowd and Zellner (1985) and Poterba and Summers (1986) use CPS reinterview data to analyze misclassification. They find that most classification errors happen on the U–N margin, and the overall effects on flows are substantial. Due to data quality problems however, reconciled reinterview data are no longer being produced. Thus, it is impossible to update their adjustment matrices, which renders the proposed mechanisms less and less reliable, particularly because they assumed that classification errors are time-invariant.

The newly resumed official worker flows are corrected for margin error, and the 1994 redesign attempted to reduce classification errors through changing the wording and order of questions, computerization and so on (see Polivka and Rothgeb (1993)). Although the post-1994 data are hopefully less prone to misclassification, measured flows are still sensitive to the errors that remain. Due to data limitations however, most studies do not address this directly. An exception is the work of Elsby et al. (2013), who consider the original Abowd and Zellner (1985) correction and an ad hoc method they call "deNUNification", whereby they eliminate transitions between unemployment and non-participation that were reversed in the next month.

2.2 First look at job finding probabilities

Figure 3 shows time series of unadjusted job finding probabilities for the 5-state partition proposed by Jones and Riddell (2006). The four non-employed states are (*i*) unemployed on temporary layoff (L), who are characterized by *waiting* to be recalled to their job rather than searching; (*ii*) unemployed-looking (U), who are active job-seekers; (*iii*) want-a-jobs (W), who report desiring work but are not searching actively, (*iv*) other non-participants (N), who claim not wanting to work. Heavy smoothing—I took 11-month moving averages—was necessary to deal with the strong seasonality in the data.[4] The underlying survey data contains the records that could have been matched successfully across four consecutive months. In addition, respondents who were retired or disabled in any month are excluded. Workers aged 16-25 and older than 55 are also excluded because their labor market transitions are likely to be driven by special factors such as school holidays and retirement.

At this point, distinguishing unemployed-looking and on temporary layoff seems absolutely warranted as the latter group find employment approximately twice as fast. One can also see that want-a-jobs are closer to unemployed-looking than to other non-

[4]The standard method in the literature is to use the up-to-date seasonal adjustment software provided by the BLS. I chose not to use the current X-13ARIMA-SEATS here because these are purely expository figures.

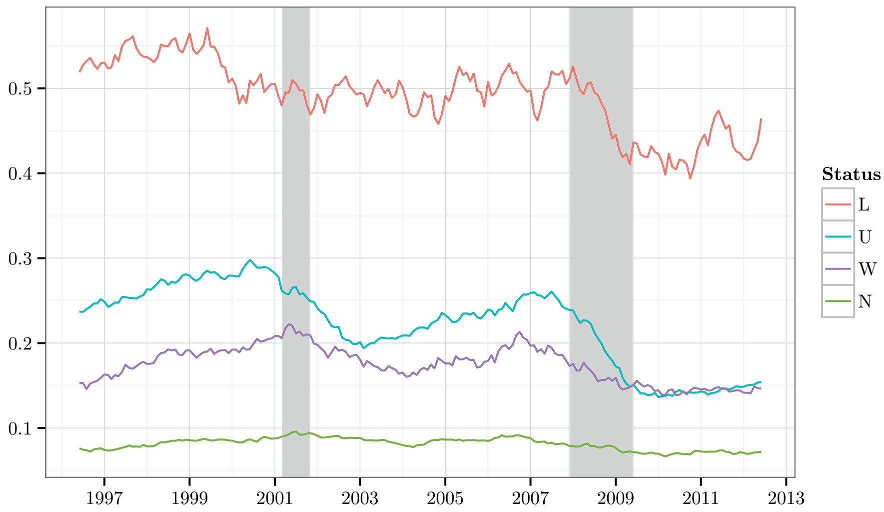

Figure 3: Job Finding Probability by Labor Force Status

Note: 11-month moving averages of the unadjusted monthly job finding probabilities of the non-employed aged 25-55. Retired and disabled workers are excluded. Matched CPS micro data from rotation groups 4 and 8.

participants. These observations are in line with Jones and Riddell (2006)'s findings for the Canadian labor market.

Notice that even these strongly smoothed series are markedly procyclical for all groups. Job finding probabilities peaked in 2006–2007, then dropped considerably as the Great Recession hit. Although to a lesser extent, the same happened in the relatively mild recession of the early 2000s. The unemployed-looking suffered the biggest decline: they were almost 40% less likely to find a job in 2009 than in 2006 as the average monthly probability fell from 24% to 14%. Want-a-jobs were affected less, their job finding probability fell by a quarter, from 19% to 15.2%. Albeit that is a considerable drop too, the gap between active job seekers and want-a-jobs seem to have closed completely. Taken at face value, this would be hard to reconcile with any standard economic model. I conjecture that it might reflect a systematic change in the way people answer the survey in recessions. The phenomenon is at least consistent, the gap between the job finding probability of the unemployed-looking and want-a-jobs narrowed in the aftermath of the 2001 recession as well. The figure is less informative on other non-participants, because their job finding probability is permanently lower than the rest's. Nevertheless, their job finding probability also fell from 9% to 7%, by more than 28%.

To motivate modeling heterogeneity within labor market states, I calculated average job finding probabilities for the cross-tabulated data in Table 1. Note that the patterns for the full sample in the first column are rather similar for the non-participants in the last column. The reason for this is that married women out of the labor force constitute by far the largest fraction of the sample. This is problematic, because non-

Table 1: Heterogeneity in Job Finding Probability (percentages)

1996–2012	Total	L	U	W	N
Sex × married					
Single women	14.1	46.9	18.5	13.7	9.7
Married women	9.0	52.9	19.6	15.1	6.5
Single men	20.2	42.2	19.5	18.9	15.3
Married men	26.6	48.5	23.6	26.5	20.6
Age					
25 to 35	12.9	46.4	21.4	16.7	7.7
36 to 45	13.2	48.6	20.8	16.8	7.8
46 to 55	13.8	47.3	17.4	16.6	8.7
Education					
Less than high school	12.9	45.9	20.1	14.6	7.3
High school	13.4	44.8	18.9	15.7	7.9
Some College	13.5	48.8	20.8	16.7	8.2
College	13.5	53.7	21.6	21.3	8.6
Region					
Midwest	13.8	46.2	19.6	17.2	8.0
Northeast	13.4	44.9	19.1	17.2	8.4
South	12.7	50.3	20.4	16.7	7.9
West	13.7	48.6	21.4	17.5	8.2
Status 3 months ago					
Employed	37.7	53.0	33.4	38.2	35.5
Unemployed on layoff	31.6	38.8	23.3	28.6	26.3
Unemployed-looking	14.2	34.1	14.8	12.1	10.5
Want a job	9.2	32.3	14.6	8.1	6.1
Non-participants	4.6	39.9	17.4	9.5	3.7

Note: Matched CPS micro data for non-employed aged 25-55 from rotation groups 4 and 8.

participant women differ markedly from other groups. First, married people find jobs quicker than singles except for women who do not want a job. Second, the gap between the job finding probabilities of unemployed and non-participants is considerably larger for women than for men. This matters as we have already seen that the gap tends to decrease in recessions. I conclude that gender and marital status are in complex interaction with labor force status. Accounting for this interaction properly would be very data-demanding, as I consider three-month histories of labor force status. On the other hand, not addressing the issue at all would distort my results as the sample is skewed towards women out of the labor force. Thus, as the focus of the thesis is on labor force history, I restrict my analysis to men.

For these prime age workers, further disaggregation by age does not seem to reveal much. However, the patterns differ across labor force states: job finding probabilities are decreasing for unemployed-looking, increasing for non-participants, and mildly parabolic for the rest. As expected, education is positively correlated with job finding,

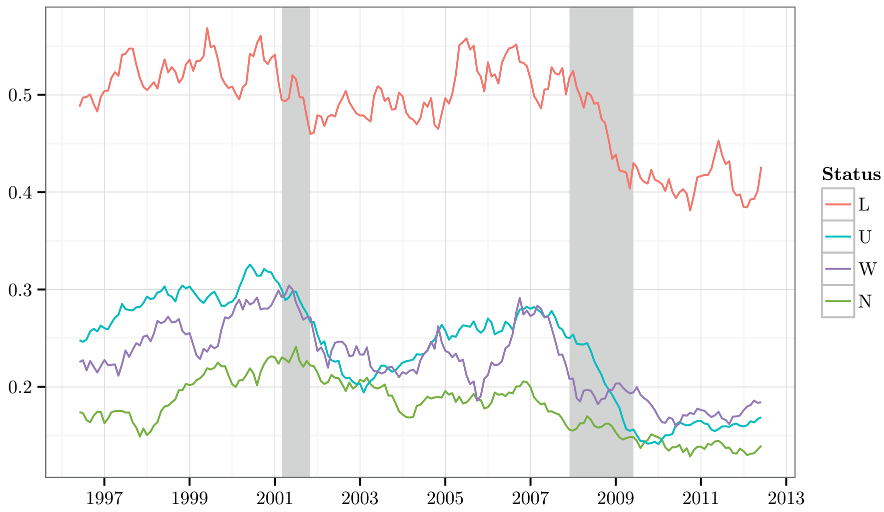

Figure 4: Job Finding Probability of Men by Labor Force Status

Note: 11-month moving averages of the unadjusted monthly job finding probabilities of the non-employed men aged 25-55. Retired and disabled workers are excluded. Includes micro data from CPS matched across four months from rotation groups 4 and 8.

although the steps are relatively small. The exception is the want-a-job category, in which college graduates have a sizable advantage. Differences between the four census regions are of similar magnitude as between educational categories. Although the effects of age, education, and region are not uniform across states either, their interactions are safer to ignore because the differences are considerably smaller and in-sample shares are more balanced than for gender.

The bottom panel attests to the profound effect of prior status. The first row shows that individuals who were employed three months earlier have very high job finding probabilities *regardless of their current status*. Although a smaller subsample, this is also true for the formerly unemployed on layoff.

The exclusion of women from the analysis reduces the sample size by 75%. This has potentially far-reaching consequences, so it is worth considering the updated version of Figure 3. The most important difference in Figure 4 is the high job finding probability of non-participants. It seems that a large fraction of prime age men start working every month, even if they said that they have not searched for or not even wanted a job. This might look suspicious, but that is what the CPS, the official employment survey of the US, tells us. On the aggregate, this is masked by the presence of stay-at-home mothers, retired and disabled, young and old persons who start working very seldom indeed.

We saw on Figure 1 that the number of N–E transitions have surpassed the number of U–E transitions in every year from 1990 onwards. As opposed to the other figures and tables in the thesis, those flows are adjusted for population weights, and are representative of the U.S. The additional insight from Figure 4 is that these high N–E flows are not just incidental transitions that become relevant only because there are many non-participants. Rather, there are people classified as out of the labor force who find jobs almost as quickly as the unemployed. Jones and Riddell (1999) argued that want-a-jobs are at least marginally attached to the labor market. In addition, I find that among prime age men, even the rest of the non-participants display considerable attachment to the labor market.

3 Models with short-term labor force history

Flows between labor market states are often modeled as a Markov chain, in which the current state contains all relevant information for predicting transition rates to subsequent states. A prime theoretical example is the DMP model, in which all unemployed and employed workers are homogeneous. Thinking in terms of Markov transitions permeates empirical work as well, see for example Shimer (2012) and Jones and Riddell (2006). These papers analyze heterogeneities in transition probabilities, but do not consider history dependence in any form.

In this section, I investigate the empirical validity of the Markov assumption by using data on short-term labor force history. In particular, I fit binomial logit models of the following form

$$\Pr(y_i = 1) = \text{logit}^{-1}(\alpha_k + X_i\beta),$$

where i is the index of individuals, and k is the index of labor market states; y_i is the indicator of finding a job in the next period, and X_i is the vector of covariates. In the first version of the model, the α_k coefficients correspond to status in the month before, while in the second they reflect history as well. I will show that paths ending in the same state correspond to significantly different job finding probabilities. Moreover, the richer model fits well the observed negative duration dependence of the unemployed, which is not targeted explicitly.

I fit the models with a general-purpose Markov Chain Monte Carlo method, the No-U-Turn Sampler. The algorithm is implemented in C++ and has user-friendly interfaces with other languages, including R. For details, see Stan Development Team (2014a,b,c). In each case, I run four Markov chains with 1000 iterations. The first 500 iterations are discarded as burn-ins, resulting in 2000 simulated sets of parameters for each model. I always check for convergence and mixing of the chains but do not discuss them in the main text. Appendix B contains the RStan outputs.

The covariates I use are age, age squared, marital status, education and region. Age, which originally takes on values between 25 and 55, is standardized by centering and dividing by two standard deviations as suggested by Gelman and Hill (2007). Education is a factor with four levels: less than high school, high school, some college, and college degree or higher. Region includes the four census regions: Northeast, Midwest, South, and West.

The choice of covariates is standard. Jones and Riddell (2006) estimate binomial logit models with the same regressors for Canadian data. The only difference is that they have a gender dummy, which I do not need, since I excluded women from the analysis. Shimer (2012) considers heterogeneity in age (7 groups), sex, race (white or nonwhite), marital status (married spouse present, spouse absent or separated, divorced or widowed, never married), nine census divisions, education (5 groups), and reason for unemployment (job loser on layoff, other job loser, job leaver, re-entrant, new entrant).

Elsby et al. (2013) takes the full interaction of gender, age (3 groups), education (4 groups), reason for unemployment (job leaver, job loser, entrant), and labor force status one year prior to the survey (E, U, N) into account. Note that the latter two authors do not fit econometric models, just apply their accounting schemes separately to the different groups.

Although Shimer (2012) has finer marital status and region partitions, he finds that heterogeneity matters little, whereas Elsby et al. (2013) can trace a large part of cyclicality in U–N transition rates to compositional changes. While taking interactions into account is potentially important as well, the authors themselves point to the inclusion of prior status as the most important difference. My goal is to elaborate on this conjecture, hence I do not partition the unemployed by reason for unemployment, because that would dramatically increase the number of possible three-month histories.

As I discussed in section 1, job finding probability over time is driven to a large extent by direct effects of aggregate fluctuations. Using several years' data without controlling for macro conditions would probably bias the results. We know that job finding probability is strongly procyclical, but we do not know whether groups with different histories are affected symmetrically. There would be numerous ways to include some indicators of macro conditions in a multilevel model. This flexibility is useful but, absent prior information, potentially dangerous.

Gelman and Hill (2007) recommends to use simple models to explore the data and inform more elaborate models. Following their advice, I start with modeling history dependence in isolation from macro conditions. To get a sense of cyclical effects, I fit the models separately for 2006 and 2009. The idea is that one year is short enough for macro conditions to change little, but has sufficiently many observations and balances out seasonality. Comparing results of a boom and a crisis year is a "non-parametric way" to learn about the relation of history dependence and business cycle and paves the way to richer models.

Table 2 contains the number of observations in the chosen years after all exclusions. In 2006, 10% and 16% of the respondents were classified as unemployed on layoff and want-a-job, respectively. The remaining three-quarters were split almost equally between unemployed-looking and other non-participants. The relevant sample for 2009 is almost twice as large. The most important difference is the dramatic increase in the number and share of respondents who claim to be unemployed. The bottom panel shows the same for a partition by history. Notice that this classification is not strictly finer, because the two subcategories of unemployment and non-participation are pooled together. Even so, there are groups with as few as 20–50 observations. The panel reveals that most of the increase in the number of currently unemployed is accounted for by the consistent job-seekers (UUU). Notice that the number of NNN remained largely the same, which corresponds to a 10 percentage point fall in the share of "permanently" out of the labor force. Lastly, EEN were the fourth largest group in 2006, but their share almost halved by 2009.

Table 2: Number of Observations

	status or history	number 2006	2009	share 2006	2009
1	L	431	919	10.62	12.49
2	U	1534	4021	37.81	54.63
3	W	654	957	16.12	13.00
4	N	1438	1463	35.44	12.49
1	UNU	95	248	2.34	3.37
2	NNN	1045	1156	25.76	15.71
3	UNN	125	186	3.08	2.53
4	NUN	93	107	2.29	1.45
5	UUN	140	267	3.45	3.63
6	NUU	126	275	3.11	3.74
7	NNU	111	182	2.74	2.47
8	UUU	711	2499	17.53	33.95
9	ENU	73	87	1.80	1.18
10	ENN	170	152	4.19	2.07
11	EUN	70	96	1.73	1.30
12	EUU	267	604	6.58	8.21
13	UEN	32	52	0.79	0.71
14	NEN	66	52	1.63	0.71
15	NEU	26	22	0.64	0.30
16	UEU	82	154	2.02	2.09
17	EEU	474	869	11.68	11.81
18	EEN	351	352	8.65	4.78
—	—	4057	7360	1	1

Note: Matched CPS micro data for non-employed aged 25-55 from rotation groups 4 and 8.

Table 3 shows summary statistics for the covariates. First, we can see that the composition of the non-employed changed little between 2006 and 2009. Thus, composition effects with respect to marital status, age, education and region are unlikely to be relevant drivers of the job finding probability over time. Second, the fraction of married workers varies substantially across labor market states. There are 25% more husbands among L than the non-employed average, which is pulled down by W and N. Third, the age distribution is remarkably symmetric, except for non-participants who are slightly younger than the others. Fourth, the non-employed with high school diploma or less were more likely to be on temporary layoff, while relatively many people with at least some college level education were out of the labor force in both years. Fifth, disproportionately many respondents were on temporary layoff in the Midwest, while relatively few in the South.

Table 3: Summary Statistics for Covariates

2006	Total	L	U	W	N
% married	44	55	46	37	41
Age					
1st quartile	30	32	32	30	29
2nd quartile	38	40	40	38	36
3rd quartile	46	47	47	47	45
Education					
% No High School	18	23	17	20	16
% High School	35	45	37	37	28
% Some College	19	17	18	17	22
% College	28	15	27	26	34
Region					
% Midwest	24	35	25	24	21
% Northeast	20	21	20	19	21
% South	29	18	31	30	30
% West	26	26	25	27	28

2009	Total	L	U	W	N
% married	47	56	49	39	40
Age					
1st quartile	30	32	31	30	28
2nd quartile	39	40	40	38	35
3rd quartile	47	48	48	46	45
Education					
% No High School	16	19	16	17	14
% High School	38	48	37	42	30
% Some College	19	15	19	17	23
% College	27	18	27	24	33
Region					
% Midwest	24	35	23	23	22
% Northeast	18	18	18	20	18
% South	30	22	31	29	33
% West	27	24	28	28	26

Note: Matched CPS micro data for non-employed aged 25-55 from rotation groups 4 and 8.

3.1 A standard model without history dependence

My benchmark model without history is a Bayesian binomial logit with flat priors

$$\begin{aligned} \Pr(y_i = 1) &= \text{logit}^{-1}(\alpha_k + X_i\beta), \\ k &\in \{L, U, W, N\}, \\ p(\alpha, \beta) &\propto 1, \end{aligned} \tag{1}$$

that is essentially one of the binomial models of Jones and Riddell (2006) with the necessary modifications for American data. Of course, this model could easily be estimated with maximum likelihood but, in anticipation of the more complex models with history, it is better to consistently use Bayesian methods throughout the thesis.

Figures 5 and 6 visualize the posterior distribution of job finding probability for 2006 and 2009 together. Figure 5 compares the α_k coefficients, hence gives a sense of job finding probability net from the effects of the other observables. In contrast, Figure 6 shows fitted probabilities on the left-hand side of (1) conditioning on all covariates. The two graphs look very similar, and the basic results can be seen from either of them. Similarity of conditional and unconditional job finding probabilities is meaningful itself, and I will return to it later.

In light of Figure 4 that pictured the raw time series, these results are not surprising. What we learn from this simple model is the magnitude of uncertainty. First, following Jones and Riddell (1999, 2006) we can ask whether the job finding probabilities out of the four states are significantly different or not. Clearly, the temporarily laid-offs are real outliers in the U.S. as well as in Canada. On average, they resume working much more quickly than any other group can find new jobs. In comparison, the three other groups are similar. After excluding women as well as old and young people, the job finding probability of the remaining non-participants (prime age men) is remarkably high. In 2006, it was around 23% for want-a-jobs and just below 20% for the rest.

The figures also reveal that the unemployed-looking were only slightly more likely to find employment than want-a-jobs in 2006. Moreover, their job finding probability plummeted by 2009 to the level of non-participants. Taking the two years together though, U, W, and N are clearly pairwise distinct states: there is basically no overlap between the point clouds. The graphs also give a sense of cyclical change: distance from the 45° line is proportional to the decline in job finding probability. What I find most striking is the dramatic decrease in the job finding probability of the unemployed-looking. The other groups seem to have weathered the crisis better in this respect, especially the want-a-jobs.

How much of this is attributable to composition effects? Based on the results of Shimer (2012) and Tables 1 and 3, we can expect that not much. First, Shimer considered all of these covariates, some of them with more categories, yet he found weak compositional effects. Second, Table 1 revealed that job finding probability varies relatively little across age, education and region categories. However, being married is

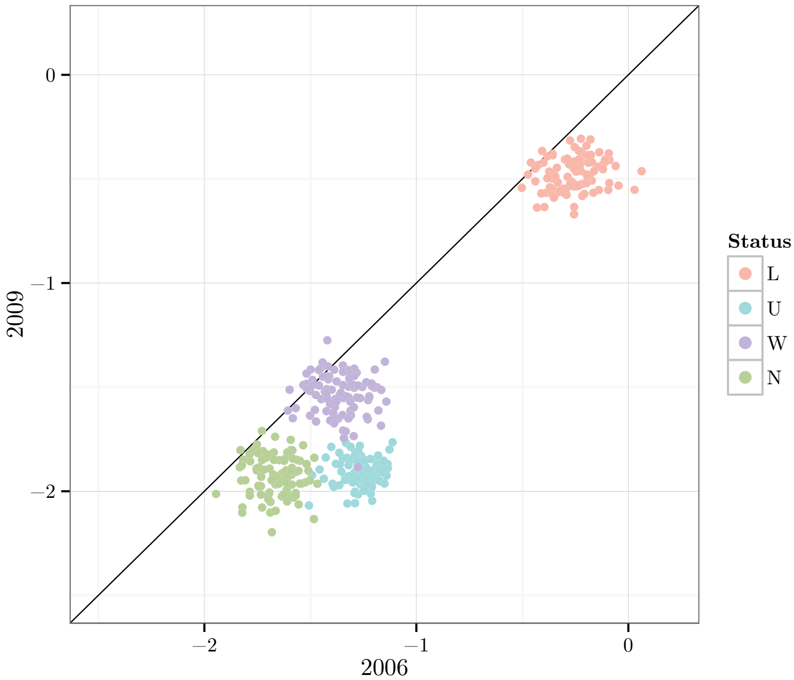

Figure 5: Coefficients in 2006 and 2009—No History

Note: A sample of 100 coefficients were drawn for each state from the total 2000 simulations.

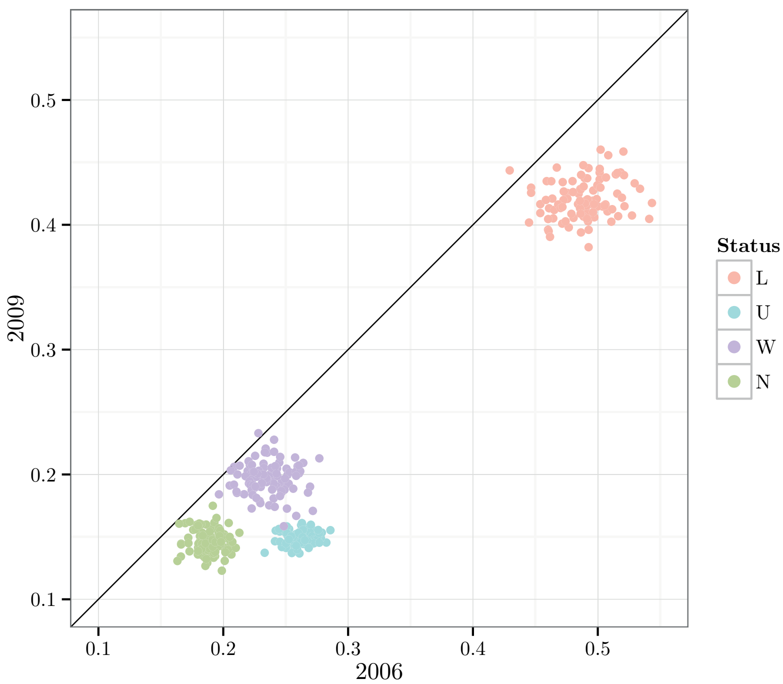

Figure 6: Fitted Job Finding Probabilities in 2006 and 2009—No History

Note: Averages of fitted job finding probabilities by labor market states
with a sample of 100 set of parameters from the total 2000 simulations.

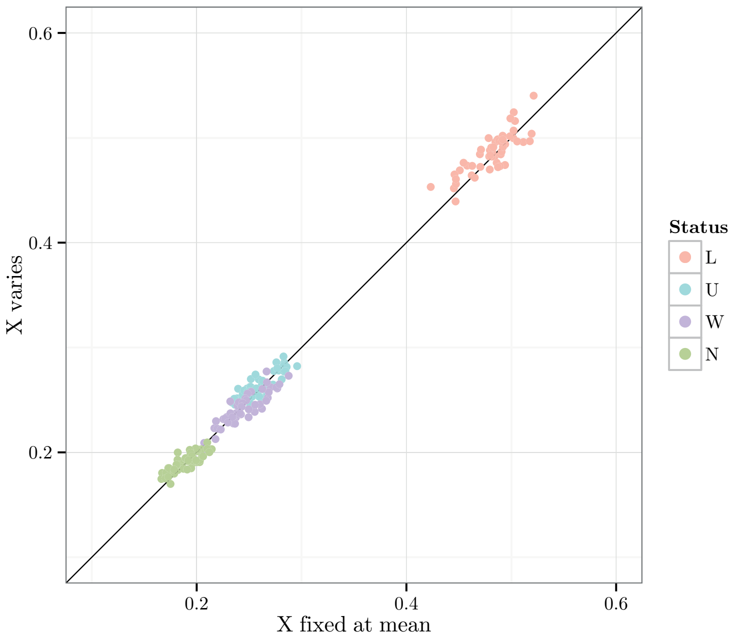

Figure 7: Coefficients vs. Job Finding Probabilities for 2006

associated with 30% higher monthly job finding probability on average. Third, Table 3 showed that the composition of the non-employed remained stable with respect to all covariates. Some insights into this question can be gained by comparing Figures 5 and 6. If the relative position of point clouds were considerably different, that would mean that the covariates in X are relevant drivers of the job finding probability over time. Although the four groups are more clearly distinct when all covariates are conditioned upon, the differences are minor.

Another way to assess the importance of within-state heterogeneity is to plot fitted probabilities on the left-hand side of (1) against fitted probabilities with covariates X fixed at their mean values. If the covariates other than labor force status played no role, all the points would line up on the 45° line. Figure 7 shows that there is some variation in job finding probabilities due to the covariates, but they are small compared to the effect of status. In conclusion, the four non-employed groups are either similar with respect to the covariates, or the covariates explain little of the job finding probability.

My probability model confirms that composition effects *within* the four states with respect to age, marital status, education, and region are small. This does not rule out however, that the average job finding probability of all non-employed had shifted due to changing population shares of L, U, W, and N. To address that in a meaningful way, I would have to re-weight the CPS sample to make it sure that it is representative of the relevant population. That is however out of the scope of my master's thesis and is left for future research.

3.2 Modeling history dependence

In full generality, all variations of three-month histories are potentially distinct states. In this subsection, I discuss a binomial logit model that allows for that possibility, while remains as close to Model (1) in every other respect as possible. In particular, I maintain the assumption that all status coefficients are independent and use the same control variables and data. Unfortunately, one year's data are insufficient for precise estimation for the 5-state partition I have used so far.[5] Therefore, I switch to the standard 3-state classification, which gives rise to eighteen different paths.[6] The model is

$$\begin{aligned} \Pr(y_i = 1) &= \text{logit}^{-1}(\alpha_{k_3 k_2 k_1} + X_i \beta), \\ k_3, k_2 &\in \{E, U, N\}, \quad k_1 \in \{U, N\}, \\ p(\alpha, \beta) &\propto 1, \end{aligned} \tag{2}$$

where k_n refers to status n months before the potential transition to employment, hence k_3 is the earliest observed status and k_1, the last, cannot be employment.

Figure 8 summarizes the posterior distribution of the estimated job finding probabilities for 2006. Significant heterogeneity with respect to history is immediately apparent. According to the model, EEU and EEN both have more than 50% chance of finding new employment. Moving down on the plot, there are all those who had recently been employed. Negative duration dependence is pronounced for unemployed and non-participants alike: job finding probability is higher for those who held jobs at $t-2$ rather than at $t-3$, though the differences are not always significant. Non-participants with no recent employment have largely the same job finding probability, which is lower than that of the unemployed. For the recently employed however, labor force history accounts for more variation than current job search behavior. I show only the estimations for 2006 but the pattern is similar for 2009. Interdecile ranges are even sharper as the 2009 sample is almost twice as large.

In their search and matching model, Kroft et al. (2014) assumed that job finding probability from unemployment is subject to causal negative duration dependence, while non-participants are re-employed at a homogeneous rate. Although my model does not identify causal relations, it demonstrates that labor force history matters at least as much for non-participants as for the unemployed. In fact, the range of fitted job finding probabilities is even larger for non-participants. Therefore, homogeneity of transition rates from non-participation is at odds with the data, even for the restricted sample of non-disabled, prime age men.

Now I turn to the comparison of boom and crisis years again. Figure 9 plots the posterior mean job finding probabilities in 2006 and 2009. We can see that there was a large dispersion in job finding probabilities in both years, especially for non-participants. Interestingly, almost all subgroups of the unemployed were hit harder

[5] With K states of which one is employment, there are $(K-1)K^2$ different three-month histories ending in non-employment.

[6] See Table 2 for the sample sizes.

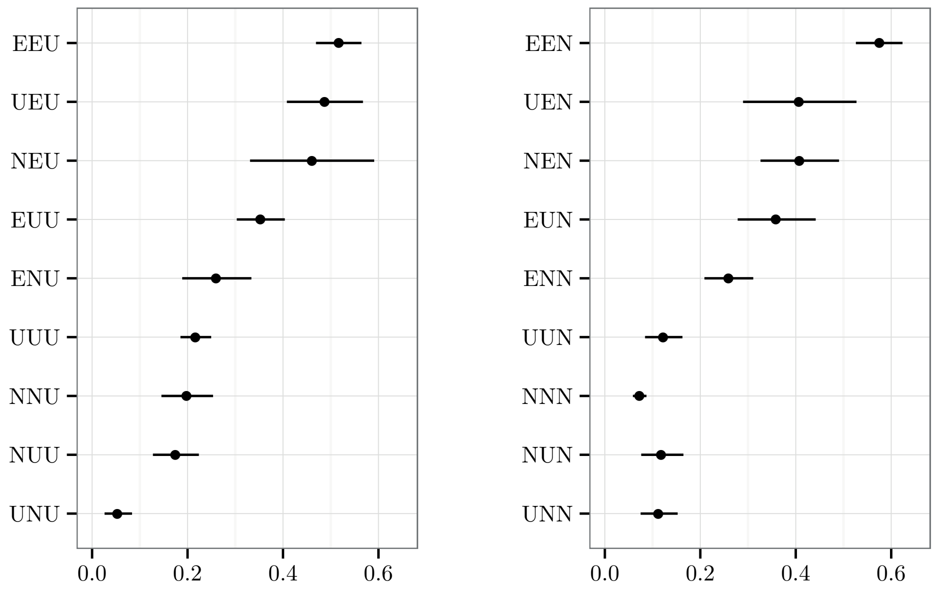

(a) unemployed

(b) non-participants

Figure 8: Posterior Uncertainty of Job Finding Probabilities in 2006—History

Dots represent means, lines interdecile ranges of the group averages across simulations.

by the crisis than non-participants. Except for UNUs, who had a surprisingly low job finding probability in 2006, which is probably an outlier. The history and no history models therefore both suggest that the job finding probability of the unemployed is more strongly procyclical than that of the non-participants.

3.3 Comparing the two models

Model (1) controls for a finer classification of the non-employed, while Model (2) includes history. If the Markov assumption were true, then controlling for history would only be useful to the extent that it picks up the effect of contemporaneous, omitted heterogeneity. In fact, many economists think that the history (or duration) dependence that we observe in the data is due to unobserved heterogeneity, see for example Hornstein (2012). Although causal effects are not identified separately in my framework, I will demonstrate that the 3-state model with history contains more information than the 5-state model without. The only potential caveat to deal with is the higher degrees of freedom of the history model.

One option is to look at out-of-sample forecasts, for which the difference in degrees of freedom is not troublesome. Table 4 shows mean squared errors (MSE) for the full samples and by the four non-employed states. The posteriors of the models fitted using 2006 and 2009 data are used to predict transition probabilities in the first halves of 2007 and 2010, respectively. The choice of half a year is somewhat arbitrary. The idea is that its macro environment should be sufficiently close to the estimation period, and

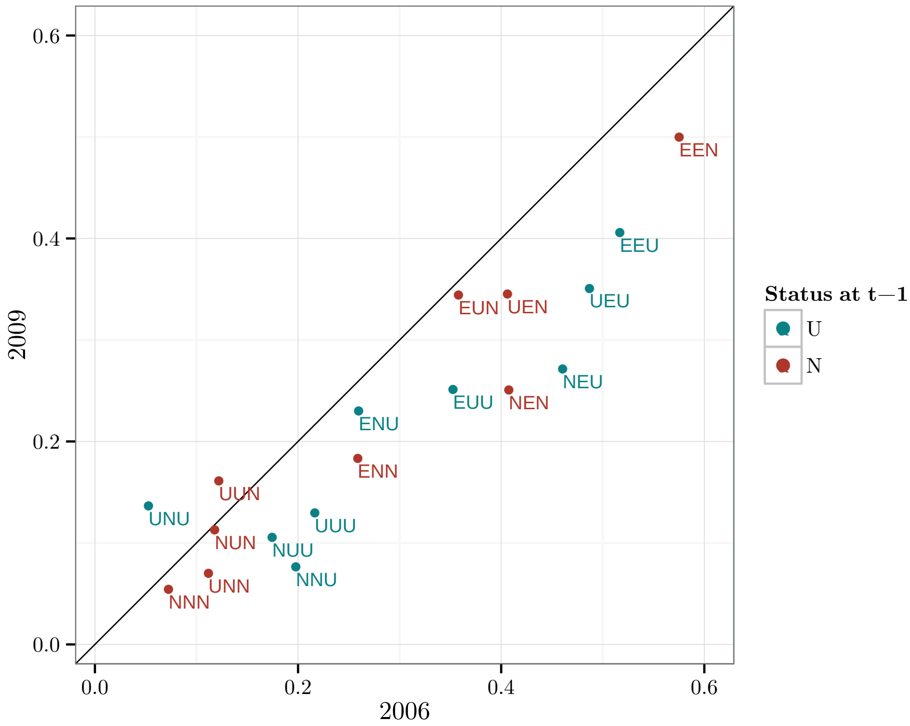

Figure 9: Fitted Job Finding Probabilities in 2006 and 2009—History

that more data is better. For this exercise, I make use of all 2000 posterior simulation, resulting in 2000 separate mean squared errors. Table 4 shows the means and standard deviations of these MSEs across the simulations.

We can see that Model (2) with history gives better predictions for both years and every group. There is one exception, temporarily laid-offs in the first half of 2010, but the difference between the two models is insignificant there. In general, the history model does considerably better for non-participants and marginally better for the unemployed. Recall however that the job finding probabilities of groups L, U, W, N were modeled explicitly in the no history model only. We have seen in section 3.1 that workers in L and U have widely disparate re-employment probabilities. Model (1) uses this information directly. Thus, it is non-trivial that the history model gives significantly better predictions for the unemployed, too.

There is another possibility to learn about model fit. The output of Bayesian models is the joint distribution of parameters given the data and the model specifics called the posterior. The posterior summarizes parameter inference and can be used to replicate the dataset via simulations. Comparison of the actual and simulated datasets can help to better understand the model fit. In Bayesian jargon, this exercise is called posterior predictive check and is commonly used in empirical work. Gelman and Hill (2007, Chapter 24) discusses the principles of predictive checking with various applications.

What aspects of the dataset would be interesting to replicate? In the introduction, I motivated modeling worker heterogeneity with the example of negative duration

Table 4: Out-of-Sample Mean Squared Errors

Model	All	L	U	W	N
			2007 Q1–Q2		
No history	0.198	0.255	0.208	0.210	0.157
	(0.001)	(0.003)	(0.001)	(0.002)	(0.001)
History	0.171	0.249	0.193	0.166	0.114
	(0.001)	(0.002)	(0.001)	(0.002)	(0.001)
			2010 Q1–Q2		
No history	0.140	0.237	0.122	0.135	0.130
	(0.004)	(0.002)	(0.000)	(0.001)	(0.001)
History	0.131	0.238	0.118	0.114	0.107
	(0.004)	(0.002)	(0.000)	(0.002)	(0.001)

dependence in job finding probability. The two panels of Figure 2 demonstrate that unemployment duration has both qualities necessary for being a relevant source of composition effects. First, job finding probability varies a lot by unemployment duration: the short-term unemployed are more than twice as likely to find employment than the long-term unemployed. Second, the population shares of short- and long-term unemployed are notably cyclical, with the ratio of long-term unemployed rising is recessions. For these reasons, it is worthwhile to compare the ability of models (1) and (2) to capture negative duration dependence. As I mentioned before, the CPS does not have duration data for non-participants. Therefore, I use the unemployed subsample.

The exact procedure for Model (1) is as follows. First, I calculate the fitted job finding probabilities

$$\widehat{p}_i = \operatorname{logit}^{-1}(\alpha_k + X_i\beta)$$

of unemployed individuals for 100 set of parameters sampled form the posterior. Second, I simulate binary U–E transitions as Bernoulli trials with success probabilities $\widehat{p}_i$ that are unique for every individual. Although unemployment duration is not used in the model, it is observed for all unemployed in the sample. Thus, it is possible to calculate the average number of transitions per units of unemployment duration. All in all, only the number of transitions is endogenous, while the covariates X_i and unemployment duration data are taken exogenously, and are the same across the 100 simulated datasets. Note that the simulated transitions reflect both parameter uncertainty and the probabilistic nature of job finding conditional on its probability. The exercise is analogous for Model (2).

Figures 10a and 10b show that the no history model implies a more or less linear decrease in job finding probability by unemployment duration. This is a pure composition effect with respect to the current state (U or L) and the covariates in X_i. Thus, while the model does capture some of the duration dependence, it falls short of explaining the sharp decrease that occurs in the first month of unemployment, and is amiss for

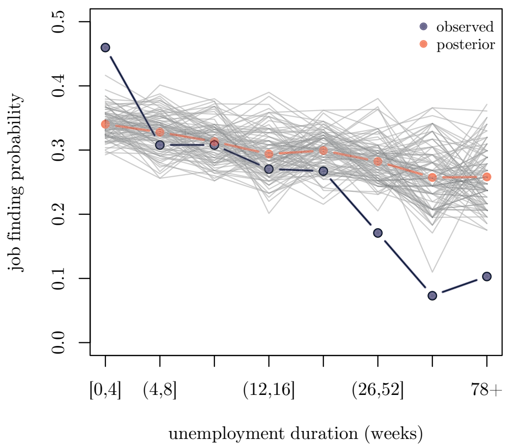

(a) no history 2006

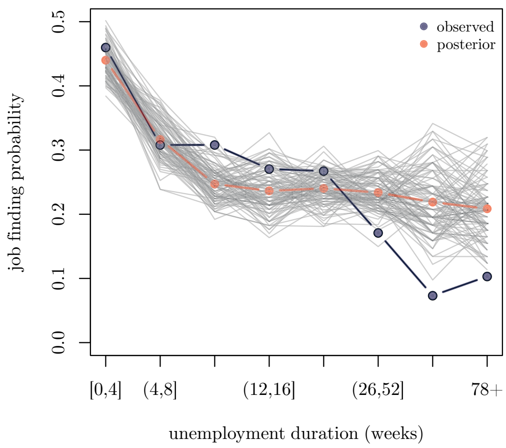

(c) history 2006

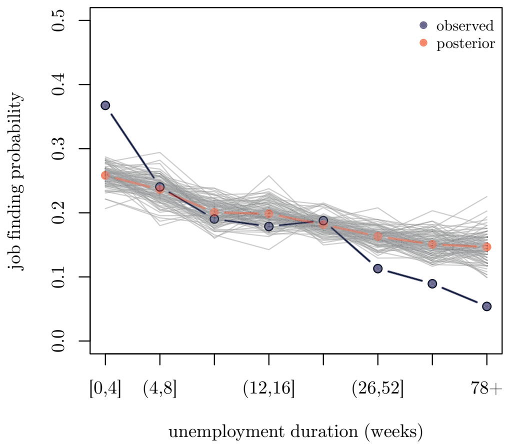

(b) no history 2009

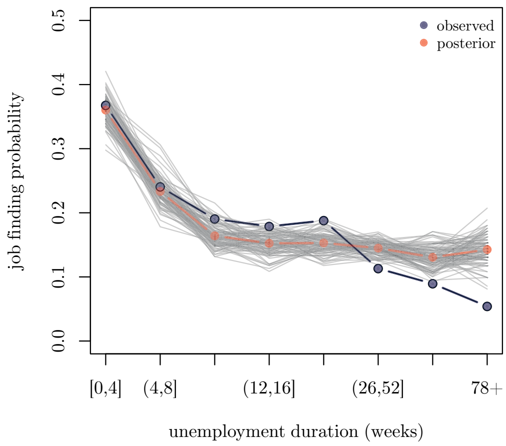

(d) history 2009

Figure 10: Posterior Predictive Check: Unemployment Duration

The blue line shows the average number of transition to employment by bins of unemployment duration as observed in the data. There is a separate bin for each month up to 6 months, then one for every half a year up to 1.5 years. Duration longer than 1.5 year are pooled together in one bin. The gray lines show the same statistics for 100 simulated datasets. The red line is the mean of these 100 simulations.

long-duration as well. In contrast, the history model fits very well for duration shorter than half a year, as can be seen from Figures 10c and 10d. Even though it does not fit well for long duration, it is significantly closer than the no history model. We can conclude that most of the observed duration dependence in job finding probabilities can be successfully accounted for by controlling for a short sequence of prior status.

Table 5: Distribution of L and W across Histories

	history	share of L	history	share of W
1	EEU	0.40	EEN	0.34
2	UEU	0.29	UEN	0.56
3	NEU	0.23	NEN	0.33
4	EUU	0.27	EUN	0.56
5	ENU	0.19	ENN	0.31
6	UUU	0.15	UUN	0.57
7	NUU	0.07	NUN	0.51
8	UNU	0.05	UNN	0.42
9	NNU	0.05	NNN	0.22

Note: Data comes from the 2006 sample.

Recall from Table 4 that inclusion of history improves model fit considerably more for non-participants than for the unemployed. Intuitively, this might be the case because history—recent employment, most importantly—is likely to be strongly correlated with being on temporary layoff, because those workers are supposed to have had a job from which they were sent home. It is less clear however, why would such a connection exist between wanting to work and duration short-term history. Table 5 provides some evidence for this intuition. Recent employment is clearly associated with a higher share of temporarily laid-offs. Although there is some regularity in the distribution of want-a-jobs as well—they typically have histories with a recent unemployment spell—that pattern is less pronounced. In addition, recent unemployment is less informative about current job finding probability than recent employment. All in all, distinguishing U and L captures a lot of heterogeneity in the job finding probability of the unemployed, but short-term history achieves the same for non-participants, as well.

4 A multilevel model of history dependence

In this section, I develop a more elaborate model to quantify the relative importance of prior labor force status at different points of time. For precise results, I need more than one year's data, thus I pool years 2002–2006. This is the five-year boom preceding the Great Recession, with relatively stable macro environment. Nevertheless, I consider including year dummies to take out fixed business cycle effects.

I keep the generality of the previous models in that there is a separate coefficient for every path. However, they are not going to be independent any more. On the one hand, this is necessary for the analysis of variance as we will see below. On the other hand, imposing a mild structure has theoretical appeal, too. To put it simply, it is a priori unlikely that NUU differ as much from UUU as from EEN. The results of the previous section confirm this intuition. Then the question is, what kind of structure is appropriate? The Markov model, which is a generally accepted approximation, says that individuals who end up being unemployed in month 3 are *the same* regardless of their former status. Taking this as a starting point, I postulate that these people are at least *more alike* then those who become non-participant in a way I formalize with the multilevel model below.

$$
\begin{aligned}
\Pr(y_i = 1) &= \operatorname{logit}^{-1}(\alpha_{k_3k_2k_1} + X_i\beta + \tau_j), \\
k_3, k_2 &\in \{E, L, U, W, N\}, \quad k_1 \in \{L, U, W, N\}, \\
j &\in \{2002, \ldots, 2006\} \\
\alpha_{k_3k_2k_1} &\sim \mathcal{N}(\mu_{k_2k_1}, \sigma_3^2), \\
\mu_{k_2k_1} &\sim \mathcal{N}(\mu_{k_1}, \sigma_2^2), \\
\mu_{k_1} &\sim \mathcal{N}(\mu, \sigma_1^2), \\
p(\alpha, \beta, \tau, \mu) &\propto 1, \\
\sigma_1, \sigma_2, \sigma_3 &\sim \text{half-Cauchy}(0, 10).
\end{aligned}
\tag{3}
$$

The first-level regression is similar to Model (2), expect for the year fixed effects τ and the use of the 5-state partition. Lines four to six describe the multilevel structure. Taking a particular example first, α_{EUU} is a realization of $\mathcal{N}(\mu_{UU}, \sigma_1^2)$ which reflects that the path EUU ends in UU. The hyper-parameter μ_{UU} is estimated analogously: it is a realization of $\mathcal{N}(\mu_U, \sigma_2^2)$ because the truncated path UU ends in U. Finally, μ_U has a normal prior. The general pattern is that the coefficients of paths with the same ending are drawn from the same distribution, whereas paths with different endings come from hyper-distributions with different means. Notice that all hyper-distributions are normal[7] and only one variance parameter is estimated per level. This is a simplifying assumption, but without it identification of history-dependent variance parameters would be very weak, because we have only 4–5 paths with the same ending. This argument brings us to the next point.

[7] I experimented with hyper-distributions with heavier tails, e.g. Student's t-distributions with low degrees of freedom, but the estimated mean parameters were robust to such changes.

In order to estimate the α coefficients precisely, we need many observations per group. That is why I chose to work with the 3-state partition in Model (2). By the same logic, hyper-parameters of a multilevel model are better identified if there are many observations. However, the observations at the hyper-level are the coefficients. Thus, it is better to have many categories for identifying the hyper-parameters. Therefore, there is a trade-off between how precisely we estimate parameters of different levels. The power of multilevel models is that they not just require but also facilitate the estimation of many parameters. In an ordinary (single-level) regression like Model (1), the parameters are independent, i.e. only observations in the same category are used to estimate the corresponding coefficient. For example, α_W is just the average job finding probability of want-a-jobs (up to the logit transformation of course). In contrast, a multilevel model implements partial pooling: Model (3) uses the information that individuals with history EEU have something in common with people in the LEU, UEU, WEU and NEU groups, and uses their average job finding probabilities as well to estimate α_{EEU}.

The amount of partial pooling is endogenous. Coefficients from the same hyper-distribution are all pulled toward their common mean, but groups with many observations or markedly different average job finding probability are affected less. On the extreme, the coefficient of a group with zero observations would be exactly the common mean. Naturally, the amount of partial pooling depends on the multilevel variance parameters, too. The higher the variance, the less information do other groups carry. Therefore, the priors on the variance parameters are in effect the prior on the amount of partial pooling. In single-level regressions, the usual choice for a non-informative prior on variance parameters is the inverse-gamma distribution. For multilevel models however, Gelman and Hill (2007) recommend using the half-Cauchy distribution. A half-Cauchy distribution with a high scale parameter is a weakly informative prior, restricting σ away from very large values, which also has better numerical properties near zero.

Having obtained reasonably precise estimates of the means of the hyper-distributions, I proceed with quantifying the contribution of labor force status at different points of time. The idea is to trace back the variations in the α coefficients to variations in $\mu_{k_2k_1}$ and μ_{k_1}. As Gelman and Hill (2007, Chapter 21) describe, variation among a batch of coefficients of a multilevel model can be summarized in two ways. The first measure is the *superpopulation* standard deviation, which captures the variability of the entire distribution the parameters are drawn from. In Model (3), these are σ_1, σ_2 and σ_3. By construction, σ_t corresponds to variations in job finding probability originating from t months before a potential transition. Importantly however, the superpopulation standard deviation is relevant for determining the uncertainty about new groups. What we need here is a description of variation among the existing categories. This second measure is the *finite-population* standard deviation, which is simply the corrected sample standard deviation of the estimated coefficients. I define the slightly modified measures of variation

$$s_1 = \sqrt{\frac{1}{K_1 - 1} \sum_{k_1} (\mu_{k_1} - \mu)^2}, \tag{4}$$

$$s_2 = \sqrt{\frac{1}{K_2 - 1} \sum_{k_1, k_2} (\mu_{k_2 k_1} - \mu_{k_1})^2}, \tag{5}$$

$$s_3 = \sqrt{\frac{1}{K_3 - 1} \sum_{k_1, k_2, k_3} (\alpha_{k_3 k_2 k_1} - \mu_{k_2 k_1})^2}, \tag{6}$$

where $K_1 = 4, K_2 = 20$ and $K_3 = 100$ are the number of different coefficients at each level. Notice that s_2 is not a proper standard deviation because it is not the common mean of all $\mu_{k_2 k_1}$'s that is subtracted. Rather, the distance form the mean conditional on the next state is taken. This adjustment is warranted, as my goal is to isolate the influence of prior status at one month, while the $\mu_{k_2 k_1}$'s correspond to two-month paths. The same applies to s_3, where only the part of the variability of $\alpha_{k_3 k_2 k_1}$'s originating in month $t - 3$ is of interest.

Table 6: Finite-Population Standard Deviations

	w/o fixed effects		with fixed effects	
	mean	s.d.	mean	s.d.
s_1	0.33	0.16	0.32	0.16
s_2	0.54	0.08	0.54	0.08
s_3	0.49	0.05	0.49	0.05

Table 6 contains the means and standard deviations of the finite-population standard deviations defined above. We can see that inclusion of year fixed effects does not affect the variation in history coefficients, most of which is due to status in month $t-2$. The earliest observed status contributes slightly less, but still more than the most recent status. Although the standard deviation of s_1 is relatively large, it is significantly lower than s_2 and s_3. Inference is more precise for s_2 and s_3 because there are more categories at those levels. Although the results might be surprising at first, their interpretation is straightforward. Ceteris paribus, more recent status is more important, but month $t - 1$ is special because employment is not a valid state in it. These results suggest that having been employed recently is more important for job finding probability than current status within non-employment.

5 Conclusion

In the master's thesis, I used a novel empirical approach to study heterogeneities in labor market flows in the U.S. Using data on short-term labor force history, I developed Bayesian logit models that capture large variations in the job finding probabilities not only of the unemployed but of the non-participants, as well. In recent years, several studies have argued that it is important to include non-participants in labor-macro models (see e.g. Elsby et al. (2013)). As it is standard in the literature, I use longitudinally matched CPS data. However, non-participants constitute a large and heterogeneous group of the working age population, and intuitively not all of them are attached to the labor market. Unfortunately, the least attached groups, such as housewifes and pensioners, are overrepresented in the CPS. Therefore, I only include non-disabled, prime age men in my sample. This way, only those are included who, at least in principle, are able to work.

One insight revealed by my analysis is that the gap between the job finding probabilities of non-participants and the unemployed is very small. First, I find that non-participants are quite likely to start working even if they claimed not wanting a job. Second, I distinguish temporarily laid-offs who constitute 9–14% of all unemployed from the unemployed-looking. The rationale behind this distinction is that those on temporary layoff are not searching actively but are waiting to be called back to their old job. Temporary laid-offs are almost twice as likely to resume working than the unemployed-looking. Once the two groups are considered separately, job finding rates of the unemployed-looking are not significantly higher than that of the want-a-jobs. Furthermore, the job finding probability of the unemployed-looking is the most sensitive to the business cycle. During the Great Recession, it fell by approximately 40%, making the unemployed-looking less likely to find employment than want-a-jobs.

Introducing short-term labor force history into the model revealed substantial heterogeneities in job finding probabilities. Most notably, recent employment is associated with high job finding probability regardless of current labor market status. To investigate the relative importance of previous states, I developed a multilevel model, and found quantitative evidence that prior status accounts for more variation in job finding probabilities than current search behavior. The history model also gives significantly better out-of-sample forecasts, especially for non-participants, and fits well the relation between job finding probability and weeks spent in unemployment. Taken together, these results suggest that negative duration dependence applies to non-participants as well, and can be captured by labor force history. This is a valuable insight, because data on history is generally available, while most employment surveys do no provide duration data for non-participants.

My master's thesis is an application of Bayesian multilevel modeling. I believe this to be a very promising approach to document macroeconomic phenomena using micro level data. Logical next steps would be the inclusion of macro variables like regional unemployment rates and vacancies, or the extension to transitions on other margins.

References

Abowd, John M., and Arnold Zellner. 1985. "Estimating Gross Labor-Force Flows." *Journal of Business and Economic Statistics*, 3(3): 254–283.

Coles, Melvyn G., and Ali Kelishomi Moghaddasi. 2014. "Do Job Destruction Shocks Matter in the Theory of Unemployment?." *Unpublished manuscript.*

Darby, Michael R., John C. Haltiwanger, and Mark W. Plant. 1985. "Unemployment Rate Dynamics and Persistent Unemployment under Rational Expectations." *American Economic Review*, 75(4): 614–37.

Darby, Michael R., John C. Haltiwanger, and Mark W. Plant. 1986. "The Ins and Outs of Unemployment: The Ins Win." Working Paper 1997, National Bureau of Economic Research.

Elsby, Michael W. L., Bart Hobijn, and Ayşegül Şahin. 2013. "On the Importance of the Participation Margin for Labor Market Fluctuations." *Unpublished manuscript.*

Frazis, Harley J., Edwin L. Robison, Thomas D. Evans, and Martha A. Duff. 2005. "Estimating Gross Flows Consistent with Stocks in the CPS." *Monthly Labor Review*, 128.

Gelman, Andrew, and Jennifer Hill. 2007. *Data Analysis Using Regression and Multilevel/Hierarchical Models.*: Cambridge University Press.

Hornstein, Andreas. 2012. "Accounting for Unemployment: The Long and Short of It." Working Paper, Federal Reserve Bank of Richmond.

Hornstein, Andreas, Per Krusell, and Giovanni L Violante. 2005. "Unemployment and Vacancy Fluctuations in the Matching Model: Inspecting the Mechanism." *FRB Richmond Economic Quarterly*, 91(3): 19–51.

Ilg, Randy. 2005. "Analyzing CPS Data Using Gross Flows." *Monthly Labor Review*, 128.

Jones, Stephen R. G., and W. Craig Riddell. 1999. "The Measurement of Unemployment: An Empirical Approach." *Econometrica*, 67(1): 147–161.

Jones, Stephen R. G., and W. Craig Riddell. 2006. "Unemployment and Nonemployment: Heterogeneities in Labor Market States." *The Review of Economics and Statistics*, 88(2): 314–323.

Kroft, Kory, Fabian Lange, Matthew J. Notowidigdo, and Lawrence F. Katz. 2014. "Long-Term Unemployment and the Great Recession: The Role of Composition, Duration Dependence, and Non-Participation." Working Paper 20273, National Bureau of Economic Research.

Krueger, Alan, Alexandre Mas, and Xiaotong Niu. 2014. "The Evolution of Rotation Group Bias: Will the Real Unemployment Rate Please Stand Up?." Working Paper 20396, National Bureau of Economic Research.

Krusell, Per, Toshihiko Mukoyama, Richard Rogerson, and Ayşegül Şahin. 2010. "Aggregate Labor Market Outcomes: The Roles of Choice and Chance." *Quantitative Economics*, 1(1): 97–127.

Krusell, Per, Toshihiko Mukoyama, Richard Rogerson, and Ayşegül Şahin. 2011. "A Three State Model of Worker Flows in General Equilibrium." *Journal of Economic Theory*, 146(3): 1107–1133.

Krusell, Per, Toshihiko Mukoyama, Richard Rogerson, and Ayşegül Şahin. 2012. "Is Labor Supply Important for Business Cycles?." Working Paper 17779, National Bureau of Economic Research.

Pissarides, Christopher. 2000. *Equilibrium Unemployment Theory, 2nd Edition.*: The MIT Press.

Polivka, Anne E., and Jennifer M. Rothgeb. 1993. "Overhauling the Current Population Survey: Redesigning the Questionnaire." *Monthly Labor Review*, 116(9): 10–28.

Poterba, James M., and Lawrence H. Summers. 1986. "Reporting Errors and Labor Market Dynamics." *Econometrica*, 54(6): 1319–1338.

Shimer, Robert. 2005. "The Cyclical Behavior of Equilibrium Unemployment and Vacancies." *American Economic Review*, 95(1): 25–49.

Shimer, Robert. 2012. "Reassessing the Ins and Outs of Unemployment." *Review of Economic Dynamics*, 15(2): 127–148.

Stan Development Team. 2014a. "RStan: the R interface to Stan, Version 2.5.0." URL: http://mc-stan.org/rstan.html.

Stan Development Team. 2014b. "Stan: A C++ Library for Probability and Sampling, Version 2.5.0." URL: http://mc-stan.org/.

Stan Development Team. 2014c. *Stan Modeling Language Users Guide and Reference Manual, Version 2.5.0.* URL: http://mc-stan.org/.

A Definitions of labor force states

In this section, I present the precise definitions of labor force states that are relevant to the thesis. See the BLS website[8] and Polivka and Rothgeb (1993).

1. **Unemployed-looking for work (U):** All those who did not have a job at all during the survey reference week, made at least one specific active effort to find a job during the prior 4 weeks, and were available for work (unless temporarily ill). An active job search method is defined as any effort that could have resulted in a job offer without any further action on the part of the job-seeker.

Active	**Passive**
Contacted:	Looked at ads
employer directly/interviewed	Attended job training programs/courses
public employment agency	Other passive
private employment agency	
friends or relatives	
school/university/employment center	
Sent out resumes/filled out applications	
Placed or answered ads	
Checked union/professional association registers	
Other active	

Table 7: Job Search Methods

2. **Unemployed on layoff (L):** All those who were not working and were waiting to be called back to a job from which they had been laid off. They need not be looking for work to be classified as unemployed.

3. **Marginally attached ($\mathbf{D} \subset \mathbf{M} \subset \mathbf{W}$):** The broadest subset of non-participants with measurable labor market attachment are those who want a job (W). Jones and Riddell (2006) calls them marginally attached.

   In official BLS terminology, one must indicate that she currently wants a job, has looked for work in the last 12 months, and is available for work in order to be counted as marginally attached to the labor force (M).

   Discouraged workers (D) report they are not currently looking for work for one of the following types of reasons: they believe no job is available to them in their line of work or area; they had previously been unable to find work; they lack the necessary schooling, training, skills, or experience; they face some form of discrimination.

[8] http://www.bls.gov/cps/cps_htgm.htm

B RStan output and convergence

Listing 1: Model (1) for 2006

```
Inference for Stan model: noh_uni_prior.
4 chains, each with iter=1000; warmup=500; thin=1;
post-warmup draws per chain=500, total post-warmup draws=2000.

        mean se_mean   sd  2.5%   25%   50%   75% 97.5% n_eff Rhat
a[1]   -0.28    0.00 0.12 -0.51 -0.36 -0.28 -0.20 -0.05   860 1.01
a[2]   -1.66    0.00 0.09 -1.83 -1.72 -1.66 -1.60 -1.47   735 1.01
a[3]   -1.26    0.00 0.09 -1.43 -1.31 -1.26 -1.20 -1.09   700 1.00
a[4]   -1.36    0.00 0.11 -1.58 -1.44 -1.36 -1.29 -1.14  1013 1.00
b_age  -0.02    0.00 0.07 -0.16 -0.07 -0.02  0.03  0.13  1733 1.00
b_age2 -0.02    0.01 0.15 -0.30 -0.11 -0.02  0.09  0.27   856 1.01
b_mar   0.46    0.00 0.08  0.30  0.40  0.46  0.51  0.61   769 1.01
e[1]   -0.05    0.00 0.06 -0.17 -0.08 -0.04 -0.01  0.06  1561 1.00
e[2]   -0.06    0.00 0.06 -0.18 -0.10 -0.06 -0.01  0.07  1569 1.00
e[3]    0.09    0.00 0.07 -0.06  0.04  0.09  0.14  0.23  1639 1.00
e[4]    0.02    0.00 0.07 -0.13 -0.03  0.02  0.06  0.15  2000 1.00
r[1]   -0.06    0.00 0.07 -0.20 -0.11 -0.06 -0.02  0.07  1584 1.00
r[2]   -0.14    0.00 0.07 -0.29 -0.19 -0.14 -0.09 -0.01  1437 1.00
r[3]    0.16    0.00 0.06  0.05  0.12  0.16  0.20  0.28  1366 1.00
r[4]    0.04    0.00 0.06 -0.07  0.00  0.04  0.09  0.17  2000 1.00

Samples were drawn using NUTS(diag_e) at Fri May 08 10:20:49 2015.
For each parameter, n_eff is a crude measure of effective sample size,
and Rhat is the potential scale reduction factor on split chains (at
convergence, Rhat=1).
```

Listing 2: Model (1) for 2009

```
Inference for Stan model: noh_uni_prior.
4 chains, each with iter=1000; warmup=500; thin=1;
post-warmup draws per chain=500, total post-warmup draws=2000.

        mean se_mean   sd  2.5%   25%   50%   75% 97.5% n_eff Rhat
a[1]   -0.48       0 0.09 -0.66 -0.54 -0.48 -0.42 -0.31  1007    1
a[2]   -1.93       0 0.09 -2.12 -2.00 -1.93 -1.87 -1.77  1287    1
a[3]   -1.91       0 0.07 -2.04 -1.95 -1.91 -1.86 -1.78   895    1
a[4]   -1.53       0 0.09 -1.71 -1.60 -1.53 -1.47 -1.35  1175    1
b_age  -0.08       0 0.06 -0.20 -0.12 -0.08 -0.03  0.05  1971    1
b_age2 -0.12       0 0.13 -0.37 -0.21 -0.12 -0.04  0.13  1257    1
b_mar   0.40       0 0.07  0.27  0.36  0.40  0.45  0.53  1090    1
e[1]   -0.10       0 0.05 -0.19 -0.13 -0.10 -0.06  0.00  1459    1
e[2]   -0.04       0 0.06 -0.15 -0.07 -0.03  0.00  0.07  1201    1
e[3]    0.09       0 0.06 -0.03  0.05  0.09  0.13  0.21  1222    1
e[4]    0.04       0 0.06 -0.08  0.00  0.04  0.08  0.16  2000    1
r[1]   -0.09       0 0.05 -0.19 -0.12 -0.09 -0.05  0.02  1315    1
r[2]   -0.03       0 0.06 -0.15 -0.07 -0.03  0.01  0.09  1333    1
r[3]    0.10       0 0.05  0.00  0.07  0.10  0.13  0.20  1556    1
r[4]    0.02       0 0.05 -0.08 -0.02  0.02  0.05  0.12  2000    1

Samples were drawn using NUTS(diag_e) at Fri May 08 11:11:20 2015.
For each parameter, n_eff is a crude measure of effective sample size,
and Rhat is the potential scale reduction factor on split chains (at
convergence, Rhat=1).
```

Listing 3: Model (2) for 2006

```
Inference for Stan model: noh_uni_prior.
4 chains, each with iter=1000; warmup=500; thin=1;
post-warmup draws per chain=500, total post-warmup draws=2000.

        mean se_mean   sd  2.5%   25%   50%   75% 97.5% n_eff Rhat
a[1]    0.08    0.00 0.13 -0.16 -0.01  0.08  0.16  0.33   920 1.00
a[2]   -0.56    0.01 0.27 -1.06 -0.74 -0.55 -0.38 -0.05  1589 1.00
a[3]   -0.53    0.01 0.38 -1.30 -0.78 -0.53 -0.29  0.21  2000 1.00
a[4]   -1.26    0.00 0.19 -1.65 -1.38 -1.25 -1.12 -0.91  1496 1.00
a[5]   -2.75    0.00 0.14 -3.03 -2.84 -2.75 -2.65 -2.50  1035 1.00
a[6]   -2.30    0.01 0.30 -2.94 -2.49 -2.29 -2.10 -1.74  2000 1.00
a[7]   -0.76    0.01 0.26 -1.30 -0.93 -0.76 -0.59 -0.25  1484 1.00
a[8]   -2.28    0.01 0.34 -2.97 -2.49 -2.26 -2.04 -1.65  2000 1.00
a[9]   -2.20    0.01 0.27 -2.76 -2.38 -2.20 -2.01 -1.67  2000 1.00
a[10]  -0.17    0.00 0.12 -0.40 -0.25 -0.17 -0.09  0.06   844 1.00
a[11]  -0.36    0.01 0.42 -1.19 -0.63 -0.36 -0.10  0.45  2000 1.00
a[12]  -0.27    0.01 0.23 -0.74 -0.43 -0.26 -0.11  0.18  1263 1.00
a[13]  -1.27    0.01 0.28 -1.83 -1.46 -1.26 -1.08 -0.73  2000 1.00
a[14]  -1.64    0.01 0.25 -2.13 -1.82 -1.63 -1.46 -1.16  1747 1.00
a[15]  -3.20    0.01 0.49 -4.31 -3.50 -3.18 -2.85 -2.33  1740 1.00
a[16]  -0.84    0.00 0.14 -1.13 -0.93 -0.83 -0.75 -0.56  2000 1.00
a[17]  -1.74    0.01 0.25 -2.25 -1.91 -1.74 -1.58 -1.27  2000 1.00
a[18]  -1.50    0.00 0.12 -1.73 -1.58 -1.50 -1.42 -1.27   833 1.00
b_age  -0.07    0.00 0.08 -0.23 -0.13 -0.07 -0.02  0.08  1763 1.00
b_age2  0.06    0.01 0.17 -0.26 -0.05  0.06  0.18  0.38   660 1.01
b_mar   0.38    0.00 0.08  0.22  0.33  0.38  0.44  0.55   842 1.00
e[1]    0.01    0.00 0.06 -0.11 -0.03  0.01  0.05  0.13  1398 1.00
e[2]   -0.06    0.00 0.07 -0.19 -0.10 -0.06 -0.01  0.07  2000 1.00
e[3]    0.05    0.00 0.08 -0.10  0.00  0.05  0.10  0.20  1525 1.00
e[4]    0.00    0.00 0.08 -0.15 -0.05  0.00  0.05  0.14  2000 1.00
r[1]   -0.05    0.00 0.07 -0.19 -0.10 -0.05 -0.01  0.08  1257 1.00
r[2]   -0.14    0.00 0.07 -0.30 -0.19 -0.14 -0.10  0.00  1473 1.00
r[3]    0.12    0.00 0.07 -0.01  0.08  0.12  0.17  0.25  1878 1.00
r[4]    0.07    0.00 0.07 -0.06  0.03  0.07  0.12  0.21  2000 1.00

Samples were drawn using NUTS(diag_e) at Mon May 11 11:41:42 2015.
For each parameter, n_eff is a crude measure of effective sample size,
and Rhat is the potential scale reduction factor on split chains (at
convergence, Rhat=1).
```

Listing 4: Model (2) for 2009

```
Inference for Stan model: noh_uni_prior.
4 chains, each with iter=1000; warmup=500; thin=1;
post-warmup draws per chain=500, total post-warmup draws=2000.

         mean se_mean   sd  2.5%   25%   50%   75% 97.5% n_eff Rhat
a[1]    -0.12    0.00 0.12 -0.35 -0.20 -0.13 -0.05  0.11  1056    1
a[2]    -1.23    0.01 0.32 -1.90 -1.44 -1.22 -1.00 -0.62  1752    1
a[3]    -0.74    0.01 0.30 -1.35 -0.94 -0.73 -0.53 -0.14  2000    1
a[4]    -1.62    0.01 0.22 -2.06 -1.76 -1.61 -1.48 -1.20  1314    1
a[5]    -2.97    0.00 0.14 -3.24 -3.06 -2.96 -2.87 -2.71  1235    1
a[6]    -2.74    0.01 0.30 -3.33 -2.94 -2.73 -2.53 -2.20  1457    1
a[7]    -0.77    0.01 0.22 -1.22 -0.92 -0.76 -0.61 -0.35  2000    1
a[8]    -2.19    0.01 0.32 -2.88 -2.39 -2.18 -1.96 -1.63  2000    1
a[9]    -1.77    0.00 0.17 -2.10 -1.88 -1.77 -1.66 -1.45  2000    1
a[10]   -0.54    0.00 0.09 -0.72 -0.60 -0.54 -0.48 -0.36   773    1
a[11]   -1.15    0.01 0.50 -2.17 -1.47 -1.13 -0.81 -0.20  2000    1
a[12]   -0.77    0.00 0.19 -1.14 -0.89 -0.77 -0.64 -0.40  2000    1
a[13]   -1.37    0.01 0.27 -1.95 -1.55 -1.37 -1.19 -0.86  2000    1
a[14]   -2.61    0.01 0.29 -3.19 -2.80 -2.59 -2.41 -2.07  1815    1
a[15]   -2.00    0.00 0.20 -2.40 -2.12 -1.98 -1.86 -1.64  1583    1
a[16]   -1.24    0.00 0.11 -1.47 -1.31 -1.24 -1.17 -1.03  1003    1
a[17]   -2.27    0.00 0.20 -2.66 -2.41 -2.27 -2.13 -1.87  1743    1
a[18]   -2.03    0.00 0.08 -2.21 -2.09 -2.03 -1.98 -1.88   723    1
b_age   -0.09    0.00 0.07 -0.22 -0.14 -0.09 -0.05  0.04  1680    1
b_age2  -0.08    0.00 0.14 -0.35 -0.17 -0.08  0.02  0.21   850    1
b_mar    0.33    0.00 0.07  0.19  0.28  0.33  0.38  0.46   897    1
e[1]    -0.05    0.00 0.05 -0.15 -0.08 -0.05 -0.01  0.05  1699    1
e[2]    -0.10    0.00 0.06 -0.21 -0.14 -0.10 -0.06  0.01  1612    1
e[3]     0.11    0.00 0.06 -0.01  0.07  0.11  0.15  0.24  1504    1
e[4]     0.03    0.00 0.06 -0.08 -0.01  0.03  0.08  0.15  2000    1
r[1]     0.01    0.00 0.06 -0.09 -0.02  0.01  0.05  0.12  1763    1
r[2]    -0.02    0.00 0.06 -0.15 -0.07 -0.02  0.02  0.10  1748    1
r[3]     0.02    0.00 0.05 -0.09 -0.02  0.02  0.05  0.12  1472    1
r[4]    -0.01    0.00 0.05 -0.12 -0.04  0.00  0.03  0.10  2000    1

Samples were drawn using NUTS(diag_e) at Mon May 11 12:24:41 2015.
For each parameter, n_eff is a crude measure of effective sample size,
and Rhat is the potential scale reduction factor on split chains (at
convergence, Rhat=1).
```

Listing 5: Model (3) with Fixed Effects

```
Inference for Stan model: year_fe.
4 chains, each with iter=1000; warmup=500; thin=1;
post-warmup draws per chain=500, total post-warmup draws=2000.


           mean se_mean   sd  2.5%   25%   50%   75% 97.5% n_eff  Rhat
mu_1      -1.29    0.02 0.49 -2.12 -1.46 -1.27 -1.08  0.42   403  1.00
mu_2[1]   -0.90    0.03 0.30 -1.50 -1.12 -0.91 -0.68  0.30   145  1.03
mu_2[2]   -1.48    0.01 0.25 -2.00 -1.64 -1.47 -1.31  1.01   959  1.00
mu_2[3]   -1.24    0.01 0.25 -1.71 -1.41 -1.26 -1.09  0.74   763  1.00
mu_2[4]   -1.41    0.01 0.25 -1.89 -1.57 -1.41 -1.24  0.92  1204  1.00
mu_3[1]   -0.16    0.01 0.26 -0.66 -0.34 -0.17  0.00  0.37   538  1.00
mu_3[2]   -0.69    0.01 0.23 -1.13 -0.84 -0.69 -0.53  0.22   604  1.01
mu_3[3]   -0.92    0.01 0.33 -1.56 -1.14 -0.91 -0.71  0.29   746  1.01
mu_3[4]   -0.91    0.02 0.27 -1.43 -1.09 -0.90 -0.73  0.38   289  1.01
mu_3[5]   -0.99    0.02 0.32 -1.57 -1.20 -0.99 -0.78  0.38   234  1.01
mu_3[6]   -0.67    0.02 0.27 -1.26 -0.83 -0.66 -0.49  0.16   239  1.02
mu_3[7]   -1.50    0.02 0.34 -2.25 -1.71 -1.49 -1.25  0.87   296  1.01
mu_3[8]   -2.13    0.01 0.24 -2.58 -2.28 -2.13 -1.97  1.65   863  1.00
mu_3[9]   -1.57    0.01 0.25 -2.06 -1.73 -1.57 -1.41  1.06   477  1.01
mu_3[10]  -1.95    0.01 0.26 -2.48 -2.12 -1.95 -1.78  1.47   572  1.00
mu_3[11]  -0.66    0.01 0.24 -1.14 -0.81 -0.66 -0.51  0.21   592  1.01
mu_3[12]  -0.98    0.01 0.26 -1.49 -1.15 -0.98 -0.82  0.46   820  1.00
mu_3[13]  -1.39    0.01 0.24 -1.85 -1.55 -1.38 -1.22  0.92   486  1.01
mu_3[14]  -1.51    0.01 0.23 -1.97 -1.65 -1.50 -1.36  1.08  1299  1.00
mu_3[15]  -1.58    0.01 0.25 -2.05 -1.74 -1.57 -1.41  1.08   693  1.01
mu_3[16]  -0.56    0.01 0.24 -1.04 -0.71 -0.56 -0.40  0.09  1100  1.00
mu_3[17]  -1.33    0.01 0.29 -1.89 -1.53 -1.34 -1.14  0.74   516  1.01
mu_3[18]  -1.66    0.01 0.26 -2.20 -1.83 -1.66 -1.48  1.16   461  1.01
mu_3[19]  -1.81    0.01 0.24 -2.29 -1.97 -1.81 -1.64  1.36   598  1.00
mu_3[20]  -2.01    0.01 0.24 -2.48 -2.17 -2.02 -1.86  1.56   783  1.00
a[1]       0.17    0.00 0.07  0.04  0.13  0.17  0.22  0.30  1046  1.00
a[2]       0.03    0.01 0.19 -0.39 -0.10  0.03  0.16  0.41   459  1.01
a[3]      -0.31    0.02 0.40 -1.09 -0.57 -0.30 -0.05  0.49   671  1.00
a[4]      -0.07    0.01 0.27 -0.59 -0.25 -0.06  0.12  0.45  1369  1.00
a[5]      -0.04    0.02 0.41 -0.93 -0.31 -0.03  0.24  0.74   464  1.00
a[6]      -0.51    0.00 0.10 -0.71 -0.58 -0.50 -0.43  0.31  1168  1.00
a[7]      -0.57    0.00 0.10 -0.77 -0.65 -0.57 -0.50  0.38   946  1.00
a[8]      -1.03    0.01 0.42 -1.86 -1.32 -1.02 -0.74  0.22  1256  1.00
a[9]      -0.45    0.01 0.26 -0.97 -0.62 -0.44 -0.27  0.04  1728  1.00
a[10]     -0.73    0.01 0.30 -1.35 -0.91 -0.72 -0.52  0.15  1334  1.00
a[11]     -0.88    0.02 0.39 -1.63 -1.13 -0.86 -0.62  0.13   446  1.01
a[12]     -0.63    0.01 0.41 -1.44 -0.92 -0.61 -0.33  0.13  1001  1.01
a[13]     -1.31    0.01 0.47 -2.25 -1.64 -1.29 -0.97  0.45  1121  1.00
a[14]     -1.10    0.02 0.51 -2.04 -1.45 -1.10 -0.74  0.12   719  1.01
a[15]     -0.72    0.01 0.47 -1.66 -1.01 -0.71 -0.41  0.17  1263  1.00
a[16]     -0.64    0.01 0.21 -1.03 -0.78 -0.65 -0.50  0.24   652  1.01
a[17]     -0.59    0.01 0.30 -1.19 -0.80 -0.59 -0.40  0.02  1173  1.00
a[18]     -1.23    0.03 0.50 -2.29 -1.54 -1.20 -0.90  0.32   274  1.01
a[19]     -0.98    0.01 0.22 -1.40 -1.13 -0.98 -0.83  0.55   489  1.01
a[20]     -1.15    0.01 0.38 -1.90 -1.42 -1.15 -0.88  0.46   888  1.00
a[21]     -0.67    0.01 0.34 -1.34 -0.89 -0.66 -0.44  0.01   750  1.00
a[22]     -1.23    0.01 0.38 -1.94 -1.49 -1.25 -0.98  0.47   675  1.00
a[23]     -1.06    0.02 0.48 -1.96 -1.38 -1.06 -0.71  0.10   779  1.00
a[24]     -0.97    0.02 0.47 -1.91 -1.27 -0.98 -0.66  0.07   673  1.00
a[25]     -1.09    0.02 0.42 -1.91 -1.38 -1.08 -0.80  0.28   439  1.00
a[26]      0.26    0.00 0.07  0.13  0.21  0.25  0.30  0.39   825  1.00
a[27]     -0.66    0.02 0.43 -1.51 -0.94 -0.66 -0.37  0.15   499  1.01
a[28]     -0.62    0.01 0.16 -0.92 -0.73 -0.62 -0.51  0.29   766  1.00
a[29]     -0.91    0.01 0.26 -1.44 -1.08 -0.92 -0.73  0.39  2000  1.00
a[30]     -0.72    0.01 0.26 -1.28 -0.89 -0.71 -0.55  0.23  1480  1.00
a[31]     -0.70    0.01 0.31 -1.33 -0.90 -0.70 -0.49  0.09   762  1.00
a[32]     -1.74    0.02 0.45 -2.66 -2.06 -1.73 -1.43  0.89   401  1.01
a[33]     -1.67    0.02 0.49 -2.64 -1.99 -1.67 -1.34  0.69   703  1.00
a[34]     -1.75    0.02 0.50 -2.78 -2.09 -1.72 -1.40  0.79   486  1.00
a[35]     -1.66    0.03 0.58 -2.91 -2.01 -1.63 -1.28  0.57   345  1.01
a[36]     -1.24    0.00 0.10 -1.45 -1.31 -1.24 -1.17  1.04  2000  1.00
a[37]     -1.75    0.01 0.39 -2.54 -2.01 -1.74 -1.48  1.00  1481  1.00
a[38]     -2.95    0.00 0.08 -3.10 -3.00 -2.95 -2.89  2.79  1294  1.00
a[39]     -2.35    0.01 0.20 -2.74 -2.48 -2.35 -2.22  1.98  1302  1.00
```

a[40]	−2.84	0.00	0.19	−3.23	−2.97	−2.84	−2.71	2.47	1645	1.00
a[41]	−0.86	0.01	0.20	−1.27	−1.00	−0.85	−0.73	0.47	1096	1.00
a[42]	−1.37	0.02	0.49	−2.41	−1.66	−1.37	−1.03	0.41	523	1.01
a[43]	−2.22	0.01	0.29	−2.80	−2.40	−2.20	−2.02	1.69	742	1.00
a[44]	−1.75	0.00	0.16	−2.08	−1.86	−1.75	−1.64	1.44	1513	1.00
a[45]	−1.66	0.01	0.24	−2.16	−1.82	−1.67	−1.50	1.20	2000	1.00
a[46]	−1.10	0.01	0.21	−1.49	−1.24	−1.10	−0.96	0.68	433	1.00
a[47]	−2.22	0.02	0.52	−3.28	−2.54	−2.20	−1.85	1.25	862	1.00
a[48]	−2.42	0.01	0.21	−2.87	−2.54	−2.41	−2.27	2.02	1020	1.00
a[49]	−1.80	0.01	0.23	−2.23	−1.96	−1.80	−1.64	1.36	2000	1.00
a[50]	−2.65	0.01	0.24	−3.15	−2.81	−2.64	−2.49	2.18	2000	1.00
a[51]	−0.51	0.00	0.05	−0.62	−0.55	−0.51	−0.47	0.40	822	1.00
a[52]	−0.50	0.01	0.26	−1.01	−0.69	−0.50	−0.33	0.02	1682	1.00
a[53]	−0.62	0.01	0.29	−1.18	−0.80	−0.62	−0.42	0.05	1705	1.00
a[54]	−0.66	0.00	0.12	−0.89	−0.74	−0.66	−0.58	0.43	1000	1.00
a[55]	−0.50	0.01	0.25	−1.00	−0.66	−0.50	−0.34	0.00	2000	1.00
a[56]	−0.84	0.00	0.18	−1.20	−0.97	−0.84	−0.72	0.50	1290	1.00
a[57]	−0.88	0.01	0.21	−1.29	−1.02	−0.88	−0.74	0.47	1058	1.00
a[58]	−0.84	0.02	0.48	−1.87	−1.15	−0.83	−0.54	0.10	829	1.00
a[59]	−1.27	0.01	0.24	−1.76	−1.43	−1.27	−1.10	0.82	2000	1.00
a[60]	−0.87	0.01	0.36	−1.60	−1.11	−0.87	−0.62	0.17	1209	1.00
a[61]	−1.15	0.01	0.22	−1.58	−1.30	−1.15	−1.00	0.71	898	1.00
a[62]	−0.92	0.01	0.39	−1.68	−1.19	−0.94	−0.66	0.13	1228	1.00
a[63]	−1.80	0.01	0.23	−2.27	−1.95	−1.80	−1.65	1.36	1124	1.00
a[64]	−1.72	0.01	0.22	−2.18	−1.85	−1.72	−1.57	1.31	287	1.01
a[65]	−1.48	0.01	0.25	−1.97	−1.64	−1.48	−1.30	0.99	542	1.01
a[66]	−1.15	0.00	0.08	−1.31	−1.20	−1.15	−1.10	1.01	634	1.00
a[67]	−1.29	0.00	0.16	−1.63	−1.40	−1.29	−1.19	0.96	1768	1.00
a[68]	−1.84	0.00	0.19	−2.22	−1.96	−1.83	−1.71	1.45	2000	1.00
a[69]	−1.75	0.00	0.05	−1.85	−1.78	−1.75	−1.71	1.66	914	1.00
a[70]	−1.77	0.00	0.12	−2.01	−1.85	−1.76	−1.68	1.55	717	1.00
a[71]	−1.02	0.00	0.15	−1.30	−1.12	−1.02	−0.92	0.73	2000	1.00
a[72]	−1.74	0.01	0.38	−2.52	−1.99	−1.73	−1.47	1.03	1235	1.00
a[73]	−1.54	0.01	0.27	−2.07	−1.72	−1.53	−1.36	1.04	589	1.01
a[74]	−1.80	0.00	0.15	−2.09	−1.91	−1.80	−1.70	1.52	2000	1.00
a[75]	−2.11	0.01	0.20	−2.51	−2.24	−2.11	−1.97	1.73	946	1.00
a[76]	−0.15	0.00	0.08	−0.29	−0.20	−0.15	−0.10	0.00	2000	1.00
a[77]	−0.03	0.01	0.32	−0.66	−0.24	−0.03	0.19	0.58	1639	1.00
a[78]	−0.21	0.01	0.29	−0.77	−0.39	−0.22	−0.03	0.40	1538	1.00
a[79]	−0.69	0.01	0.22	−1.13	−0.85	−0.68	−0.53	0.28	933	1.00
a[80]	−1.04	0.01	0.24	−1.53	−1.20	−1.04	−0.87	0.55	2000	1.00
a[81]	−0.96	0.01	0.28	−1.50	−1.16	−0.96	−0.77	0.40	666	1.01
a[82]	−1.13	0.01	0.33	−1.81	−1.34	−1.13	−0.91	0.52	1258	1.00
a[83]	−1.71	0.02	0.50	−2.78	−2.01	−1.68	−1.37	0.80	427	1.01
a[84]	−1.28	0.01	0.43	−2.14	−1.56	−1.29	−0.98	0.40	1239	1.00
a[85]	−1.52	0.02	0.48	−2.52	−1.82	−1.51	−1.18	0.66	978	1.00
a[86]	−1.34	0.00	0.22	−1.78	−1.49	−1.34	−1.18	0.90	2000	1.00
a[87]	−1.30	0.02	0.49	−2.22	−1.61	−1.33	−1.00	0.29	556	1.00
a[88]	−1.95	0.01	0.16	−2.26	−2.06	−1.95	−1.83	1.63	951	1.00
a[89]	−1.95	0.01	0.29	−2.50	−2.13	−1.93	−1.75	1.40	637	1.00
a[90]	−1.88	0.01	0.21	−2.30	−2.01	−1.88	−1.73	1.50	607	1.01
a[91]	−1.22	0.01	0.19	−1.57	−1.35	−1.21	−1.09	0.86	675	1.00
a[92]	−1.97	0.01	0.39	−2.78	−2.23	−1.95	−1.70	1.22	2000	1.00
a[93]	−2.01	0.01	0.33	−2.69	−2.23	−2.00	−1.78	1.40	1112	1.00
a[94]	−2.17	0.00	0.15	−2.48	−2.28	−2.17	−2.07	1.88	2000	1.00
a[95]	−1.98	0.01	0.21	−2.42	−2.12	−1.98	−1.84	1.60	413	1.00
a[96]	−1.13	0.00	0.19	−1.50	−1.25	−1.13	−1.00	0.76	2000	1.00
a[97]	−2.28	0.01	0.46	−3.22	−2.57	−2.26	−1.96	1.44	1632	1.00
a[98]	−2.57	0.01	0.25	−3.05	−2.74	−2.57	−2.41	2.09	573	1.01
a[99]	−2.18	0.00	0.20	−2.57	−2.31	−2.17	−2.04	1.83	2000	1.01
a[100]	−2.41	0.00	0.16	−2.74	−2.52	−2.41	−2.29	2.10	1492	1.00
b_age	−0.12	0.00	0.03	−0.18	−0.14	−0.12	−0.10	0.06	2000	1.00
b_age2	−0.02	0.00	0.07	−0.16	−0.07	−0.02	0.02	0.11	631	1.00
b_mar	0.26	0.00	0.03	0.20	0.24	0.26	0.28	0.33	688	1.00
sigma_1	0.65	0.04	0.69	0.07	0.26	0.46	0.78	2.52	337	1.02
sigma_2	0.58	0.01	0.13	0.37	0.49	0.56	0.66	0.90	626	1.00
sigma_3	0.50	0.00	0.06	0.40	0.46	0.49	0.54	0.62	284	1.01
e[1]	−0.03	0.00	0.03	−0.08	−0.04	−0.02	−0.01	0.02	843	1.00
e[2]	−0.02	0.00	0.03	−0.07	−0.03	−0.02	0.00	0.04	1390	1.00
e[3]	0.02	0.00	0.03	−0.04	0.00	0.02	0.04	0.08	966	1.00
e[4]	0.02	0.00	0.03	−0.04	0.00	0.02	0.04	0.08	2000	1.00
r[1]	−0.06	0.00	0.03	−0.12	−0.08	−0.06	−0.05	0.01	1577	1.00

```
r[2]      -0.08     0.00  0.03  -0.14  -0.10  -0.08  -0.06  0.02   582  1.01
r[3]       0.05     0.00  0.03   0.00   0.03   0.05   0.07  0.10  1533  1.00
r[4]       0.09     0.00  0.03   0.04   0.07   0.09   0.11  0.15  2000  1.00
tau[1]    -0.05     0.00  0.03  -0.12  -0.07  -0.05  -0.03  0.01  1191  1.00
tau[2]    -0.04     0.00  0.03  -0.10  -0.07  -0.04  -0.02  0.02   604  1.01
tau[3]    -0.03     0.00  0.03  -0.10  -0.06  -0.03  -0.01  0.03   838  1.00
tau[4]     0.05     0.00  0.03  -0.02   0.02   0.05   0.07  0.11  1482  1.00
tau[5]     0.08     0.00  0.03   0.02   0.06   0.08   0.11  0.15  2000  1.00

Samples were drawn using NUTS(diag_e) at Sun May 17 18:39:34 2015.
For each parameter, n_eff is a crude measure of effective sample size,
and Rhat is the potential scale reduction factor on split chains (at
convergence, Rhat=1).
```