Essays on Labour Market Search

Jingcheng Fu

Supervisors: Prof. Richard Upward
Prof. Martin Sefton

School of Economics
University of Nottingham

This dissertation is submitted for the degree of
Doctor of Philosophy

January 2018

Acknowledgements

I would like to express my deepest gratitude to my supervisors, Richard Upward and Martin Sefton, for their excellent guidance and assistance in the course of my PhD. Their incessant encouragement and support have been a true inspiration to me in these years.

I am also very grateful to Martyn Andrews with whom I have collaborated on one of the essays included in this thesis. I have learned a lot from the experience which I really enjoy and appreciate.

I would also like to thank the members of the Centre for Decision Research and Experimental Economics (CeDEx) for their many useful suggestions at various stages of this research. In particular I would like to thank Daniele Nosenzo for the comments and suggestions to Chapter 2 of this work.

My gratitude also extends to participants to the numerous seminars and conferences where I have presented my work, for their helpful comments.

I am grateful to the Economic and Social Research Council (ESRC) and to the Nottingham School of Economics for their financial support during the period of my study, and to CeDEx for providing the resources and facilities for running experiments.

I thank my great office mates in B51 for their friendship and support.

A special acknowledgment goes to my friends and professors in Hong Kong Baptist University who enlightened me with the first glance of research, encouraged me and offered a lot of support.

Last but not least, I thank Hongfei and our parents for their continued love and understanding.

Abstract

This thesis contains three studies on the topic of labour market search. Chapter 1 provides an overview of the studies.

Chapter 2 reports an experimental study which examines how social comparisons affect behavior in a sequential search task. In a control treatment subjects search in isolation, while in two other treatments subjects get feedback on the search decisions and outcomes of a partner subject. The average level and rate of decline of reservation wages are similar across treatments. Nevertheless, subjects who are able to make social comparisons search differently from those who search in isolation. Within a search task we observe a reference wage effect: when a partner exits, the subject chooses a new reservation wage which is increasing in partner income. We also observe a social learning effect: between search tasks, subjects who have been paired with a more patient and successful partner increase their reservation wages in the next task.

Chapter 3 reports a study in which we provide the first microeconometric estimates of the hazards to matching on *both* sides of a labour market, decomposed into two constituent parts. Namely, (i) the rate at which job-seekers and vacancies contact each other (i.e. having interviews), and (ii) the probability that a contact results in a match. To do this, we use unique data which contains information on job-seekers, vacancies, interviews and interview outcomes. We use a specification which addresses the problems of the temporal aggregation bias and spatial spillovers highlighted by the two-sided estimates. Our estimates suggest that market tightness affects the matching rates mainly

through affecting the meeting rates. In both the raw data and the estimates, we find the decline in the matching hazard is driven by the decline in the contact hazard, and not by a fall in the matching probability. And we also report the effects of various characteristics on matching decomposed into the effects on meeting and matching probability.

Using the same data as Chapter 3, Chapter 4 provides further evidence on the mechanism by which job-seekers and vacancies decide whom to contact during their search. Since the data features an environment where both sides of the market have access to a database (or marketplace) of potential partners, a natural model of search is one of stock-flow matching, and we show that the predictions of this model outperform those of a simple random matching model. Our descriptive and econometric evidence shows that it is the inflow rate of new agents, rather than the total stock of agents, which determines the contact rates of existing agents, consistent with the predictions of the stock-flow model.

Chapter 5 summarizes the findings of this dissertation and concludes.

Table of contents

List of figures viii

List of tables ix

1 Introduction 1

2 Social comparisons in job search: Experimental evidence 7

2.1 Introduction . . . 8
2.2 Literature . . . 13
2.2.1 Search Experiments . . . 13
2.2.2 Social Comparison Effects . . . 15
2.3 Experimental Design . . . 17
2.3.1 The basic search task . . . 17
2.3.2 Treatments . . . 18
2.3.3 Discussion of the design . . . 19
2.4 Results . . . 23
2.4.1 The impact of social feedback within sequences . . . 27
2.4.2 Impact of social feedback between sequences . . . 32
2.5 Conclusions . . . 36
Appendix 2.A Instructions . . . 40
Appendix 2.B Theoretical Analysis . . . 47
2.B.1 Basic setup . . . 47
2.B.2 Distributional preferences . . . 48
Appendix 2.C Additional results . . . 56

3 Meeting and matching: Decomposing the matching function 57

3.1 Introduction . . . 58
3.2 Literature . . . 61

3.2.1 Estimates of the matching function and its components 61
3.2.2 Duration dependence of hazard rates 62
3.2.3 Effects of agent characteristics on hazard rates 64
3.3 Econometric Methods 66
3.3.1 Decomposing job-seeker and vacancy hazards 66
3.3.2 Count data and binary choice estimation 69
3.3.3 Temporal aggregation bias 73
3.3.4 Geographic spill-overs 76
3.4 Data 79
3.4.1 Creating estimation panels 81
3.4.2 Sample statistics 83
3.4.3 Stocks and inflows 84
3.4.4 Defining local districts and regions 87
3.5 Results 88
3.5.1 Comparison of specifications 88
3.5.2 Decomposition of the matching function 91
3.5.3 Decomposition of hazard rates 94
3.5.4 Decomposition of the effect of characteristics 97
3.6 Conclusions 100
Appendix 3.A Temporal aggregation bias: simulated estimations 103
Appendix 3.B Map of Lancashire 105
Appendix 3.C Additional descriptive statistics 106
Appendix 3.D Additional regression results 109

4 Evidence on stock-flow search **110**
4.1 Introduction 111
4.2 Previous literature 114
4.2.1 Tests of stock-flow matching model 114
4.2.2 Evidence from search activities 115
4.3 Empirical Framework 116
4.4 Data 121
4.5 Descriptive Evidence 124
4.5.1 Duration dependence 125
4.5.2 Who contacts whom? 127
4.6 Econometric Evidence 129
4.6.1 Random matching model 129

4.6.2 Stock-flow matching model 131
4.6.3 Robustness 134
4.7 Conclusions 137
Appendix 4.A Additional descriptive statistics 140
Appendix 4.B Additional regression results 141

5 Conclusions **143**

References **147**

List of figures

2.1 Optimal reservation wage after partner exits 21
2.2 Reservation wage by search length 26
2.3 Correlation between changes in reservation wage and partner's offer . . 29
2.4 Period 1 reservation wages by treatment and sequence 33
2.5 Impact of partner's outcome in the previous sequence 35
2.A1 Screenshots for period feedback 44
2.B1 Equilibrium initial reservation wage 54

3.1 Organisation of the Matching Service data 82
3.2 Aggregate monthly stocks, inflows, contacts and matches 86
3.3 Comparison of raw and conditional hazard rates 95
3.4 Baseline hazards for meetings 97
3.B1 Map of Lancashire 105
3.C1 Monthly time-series of stocks and inflows by district and region 108

4.1 Total and new stocks of job-seekers and vacancies 124
4.2 Contacts per month of job-seeker and vacancy search. 125
4.3 Contacts per month of search by unemployment spell length 126
4.4 Features of contacted vacancies over unemployment duration 129
4.A1 Durations of contacted agents 140

List of tables

2.1 Treatment Design 20
2.2 Observable characteristics of subjects across treatments 24
2.3 Overview of reservation wages across treatments 25
2.4 Estimates of change in reservation wage between periods 31
2.B1 Quantifying the implications for search outcomes 55
2.C1 Alternative specifications for change in reservation wage between periods 56
2.C2 Estimates of sequence trend 56

3.1 Summary statistics of job-seeker and vacancy panels 84
3.2 Estimates of matching function using various specifications 90
3.3 Estimates of hazard rate, contact rate and matching probability 92
3.4 Estimated effect of covariates 99
3.A1 Simulated contact function estimation 104
3.C1 Summary of raw data: hazards, contacts and matching probabilities . . 106
3.C2 Observable characteristics by results of job-seeker search spell 107
3.D1 Raw coefficients and estimates of spill-over parameters 109

4.1 Summary statistics of monthly contacts 123
4.2 Who contacts whom? 128
4.3 Random matching estimates of monthly contacts 130
4.4 Random matching estimates separating old and new agents 131
4.5 Stock-flow estimates of monthly contacts 133
4.6 Stock-flow estimates of monthly contacts for first spells only 136
4.7 Stock-flow estimates of monthly contacts by sub-market 137
4.A1 Who contacts whom? Job-seekers with stable occupational choices . . . 140
4.B1 Stock-flow estimates of monthly contacts, fixed-effects estimates 141
4.B2 Stock-flow estimates of monthly contacts for various window periods . 142

Chapter 1

Introduction

This thesis is a collection of three papers (Chapters 2, 3 and 4) that contribute to the understanding of the labour market search. The thesis uses methods of experimental economics as well as draws on field data. Both laboratory and naturally-occurring data are analysed using advanced econometric methods.

Chapter 2 reports an experiment designed to examine social comparison effects in job search. The research question is inspired by the social psychology literature which shows that social comparison has a powerful influence on the way humans feel and behave (Bandura and Walters, 1977; Festinger, 1954). In addition, there is substantial field evidence that search outcomes are correlated within groups of individuals who interact (e.g. Marmaros and Sacerdote, 2002). We suggest part of such correlations may be due to social comparison effects rather than network effects that are commonly discussed in the labour economics literature. In our experiment we exclude network effects by design but allow for the possibility of social comparison effects.

We focus on two reasons that social comparisons may affect job search. One is that job-seekers may have *distributional preferences* since there is extensive evidence that interpersonal comparisons can influence individuals' sense of well-being and hence labour market outcomes (e.g. Akerlof and Yellen, 1990; Frank, 1984, 1985; Loewenstein et al., 1989). A second reason stems from *learning*, as observed in other studies which

incorporate social interactions in the decision under risk (e.g. Bursztyn et al., 2014; Offerman and Schotter, 2009).

We use a sequential search task based on McCall (1970) whereby searchers receive wage offers from a known distribution and decide whether to accept the offer, receiving this offer in each remaining period, or reject the offer and wait for another offer. The model is implemented using a random termination method, whereby subjects search with an uncertain time horizon and the continuation probability plays the role of the discount factor. Previous experiments have studied similar models (Braunstein and Schotter, 1982; Brown et al., 2011; Hey, 1982, 1987; Kogut, 1990; Oaxaca and Cox, 1989, 1992; Schotter and Braunstein, 1981; Schunk, 2009; Schunk and Winter, 2009; Sonnemans, 1998), but all of these focus on individual decisions made in isolation.

In our experiment we conduct a control treatment where subjects search in isolation, and two other treatments where subjects are paired and complete independent search tasks, and are given search feedback on their partners' decisions and outcomes.

In the control treatment we replicate the findings from the previous literature using similar search tasks: subjects tend to use a reservation wage which is too low, and decrease reservation wages over the course of search. In the social comparison treatments, we observe similar levels and rates of decline of reservation wages as the control treatment, but we also observe that subjects systematically respond to feedback about their partner's search outcomes. Within a search task we observe a reference wage effect: when a partner exits, the subject chooses a new reservation wage which is increasing in partner income. Our theoretical analysis shows that this is qualitatively consistent with the prediction of distributional preferences modeled by Fehr and Schmidt (1999) when applied to our search task. We also observe a social learning effect: between search tasks, subjects who have been paired with a more patient and successful partner increase

their reservation wages in the next task. These findings suggest that social comparisons have a real impact on search decisions.

Chapter 3 provides a microeconometric analysis of the meeting and matching process in a well-defined labour market. Our data records the search durations for individual job-seekers and vacancies, as well as any meetings (i.e. interviews) and matches (i.e. hires) that occur to them though a Matching Service. This two-sided micro data sheds light on recurring issues in the empirical matching literature. It is well-known that data recorded at discrete frequency (e.g. monthly) may lead to biased estimates of the parameters of the matching process which is instantaneous (Berman, 1997; Burdett et al., 1994; Coles and Petrongolo, 2008; Gregg and Petrongolo, 2005). It has also been highlighted that spatial aggregation can lead to biases: Coles and Smith (1996) discuss the problem with aggregating separate labour markets as a single market, and other works (Burda and Profit, 1996; Burgess and Profit, 2001; Patacchini and Zenou, 2007; Petrongolo and Manning, 2017) demonstrate the importance of accounting for interdependence between neighbouring labour markets. Uniquely, we show that temporal aggregation bias and cross-district search mean that a naive estimate of the matching function produces biased estimates, which are divergent when this is done from both the job-seeker's and the vacancy's perspectives. We also show that using a specification which incorporates suggestions of both strands of the literature causes the estimates of the parameters of the matching function to move much closer together.

We decompose the matching rates into the meeting rates and matching probability. Through the lens of this decomposition, we are able to investigate (i) the mechanisms by which the stocks of job-seekers and the vacancies affect the matching process, (ii) the causes of the declining hazard rate over the search duration, and (iii) the effects of job-seeker and vacancy characteristics. These questions are central to the search and matching literature. Previous studies tend to rely on structural assumptions to facilitate

similar decompositions because the meetings are not directly observed (e.g. Frijters and van der Klaauw, 2006; García-Pérez, 2006; Lollivier and Rioux, 2010; Petrongolo and Pissarides, 2006), or estimate the meeting and matching functions at aggregate level (Berman, 1997). Our agent-level data allows us to address the research questions more directly and accurately.

Our decomposition leads to three key findings. First, from the estimates of the matching and meeting functions, we find that constant returns to scale are not rejected for the matching function, and that labour market tightness mainly operates via the meeting decision, not through the matching probability. Second, we find that the meeting hazard exhibits a strong negative duration dependence for both job-seekers and vacancies after controlling for observed characteristics and unobserved heterogeneity, and no strong duration effect is found in the matching probability. Third, in terms of characteristics, the effects on meeting rates and matching probabilities are often in the opposite direction. For example, males have a lower meeting rate but a higher matching probability, and the combined effect is a higher matching rate than females, and job-seekers with a high education qualification have a higher matching rate, which arises from much a higher contact rate and a lower matching probability. This is not always the case, as ethnic minorities have a lower meeting rate and a lower matching probability, which result into a much lower matching rate.

These findings have important policy implications. A lot of labour market policies are centered on the issue of reducing long-term unemployment. Our findings suggest that the state should not try to force the reservation wages down (e.g. through cutting unemployment benefits). The duration pattern of matching probabilities (and hence the reservation wages) is close to constant whereas the meeting hazards fall immediately after the first month — if the reservation wage were to fall, it needs to fall by a large amount (yet still might not be adequate) to make the matching hazard constant. A related

question is whether increasing the contacts between job-seekers and employers (e.g. through forcing the benefit claimants to have a minimum number of contacts per week with potential employers) increases unemployment outflows. This will depend on the cause of the fall in the contact hazards, or more specifically, whether contact hazards fall due to lack of potentially suitable vacancies in the market.

Partly motivated by this unresolved question, Chapter 4 reports another microeconometric study which investigates the pattern of interviews between job-seekers and vacancies. The aggregate matching function introduced by Pissarides (1979) describes the random search process and is a widely-adopted tool for modelling labour market frictions. The stock-flow matching model introduced by Coles and Muthoo (1998) and Coles and Smith (1998) makes some fundamentally different assumptions about the matching process, and gives rise to a different matching function. They assume that screening potential opportunities on the other side of the market can be achieved in a relatively short period of time. Some agents remain unmatched, because agents are heterogeneous and a suitable partner does not readily exist in the market. The stock-flow model predicts that agents who remain in the market after the initial search period (the "stocks") only match with new agents (the "flows"), and therefore the exit rates of job-seekers and vacancies fall once the initial search period is over.

Tests of this model on match outflows are provided by a number of studies, using aggregate time series data (Coles and Smith, 1998; Ebrahimy and Shimer, 2010; Shimer, 2007), market-level data (Coles and Petrongolo, 2008; Gregg and Petrongolo, 2005; Sasaki, 2008; Soininen, 2007) and data of individual unemployment/vacancy spells (Alvarez de Toledo et al., 2011; Andrews et al., 2013; Kuo and Smith, 2009). In general they conclude that the stock-flow matching model outperforms the random matching model. All of these studies use matches (or hires) as their dependent variable, but the stock-flow model is actually a hypothesis about the contact process. For example,

discrimination and human capital depreciation can also lead to declining matching hazards even if the meeting process is random. Therefore a more direct test of the model requires data on interviews.

We provide new evidence on the the matching mechanism by testing implications of the stock-flow matching model in the patterns of interviews. First, we confirm the that meeting hazards decline with duration (consistent with our findings from Chapter 2). Second, we find that, as an agent spends longer in the market, the average duration of the partners they meet for interviews tend to decline. And our econometric estimates suggest that it is the inflow rate of new agents, rather than the total stock of agents, that determines the meeting rates of the old agents who have spent more than four weeks in the market. These findings are robust to several specifications, and are consistent with the predictions of the stock-flow matching model. This provides new evidence that the idea of stock-flow matching is consistent with evidence on contacts, rather than just on matches. This also implies that policies to increase the contact rates of the long-term unemployed might have limited effectiveness on increasing their job-finding rate.

Chapter 5 concludes by summarizing the results of Chapters 2 to 4, pointing out their limitations and suggesting potential future research.

Chapter 2

Social comparisons in job search: Experimental evidence

[*] This chapter is based on a joint work with Martin Sefton and Richard Upward.

2.1 Introduction

Job search is typically not undertaken in isolation. Instead, job-seekers are aware of the job search activity of friends, colleagues and acquaintances, and information about others' job search successes and failures is often available. The social nature of search introduces scope for several factors that are usually ignored in economists' models of job search. In particular, a considerable body of evidence from social psychology shows that social comparisons have a powerful influence on the way humans feel and behave (Festinger, 1954). People care about their status and evaluate this by comparing themselves with others, and people who are uncertain about what decision to make are influenced by the observed behaviour of others (Bandura and Walters, 1977). There is indirect evidence that such social comparisons affect job search. Using survey data, Clark (2003) finds that the unemployment rate in the social reference group is strongly positively correlated with the self-reported well-being of the unemployed. Thus, apparently, the unemployed feel better when they see that others around them are also unemployed. If this reduces the motivation to exit unemployment in areas with high unemployment it may contribute to regional variation in unemployment rates. More generally, social comparison effects may contribute to findings, such as those of Marmaros and Sacerdote (2002), that search outcomes are correlated within groups of individuals who interact.

However, it is difficult to attribute correlated outcomes in field data to social comparison effects because there are other ways that interaction within a group can affect search outcomes. As noted in several theoretical and empirical studies of labour markets, social interaction can also affect job search by influencing a job-seeker's available options and information about those options. For instance, graduates from the same institution may share job news and make use of alumni contacts which help them find a job. These types of *network effects*, which work through changes to the offer process and which therefore have direct effects on the expected costs and benefits of taking an offer or waiting, can

be viewed through the lens of standard economic theory in which job-seekers weigh up the expected costs and benefits of search.[1]

In this paper we focus on the case where one job-seeker's offer process is independent of another's. In this case, observing another job-seeker's decisions or outcomes may still influence behaviour due to *social comparison effects*. To isolate social comparison effects from network effects we use laboratory methods, where interaction within a group can be tightly controlled, to eliminate the possibility of network effects.

Social comparison effects on job search can arise for a variety of reasons. One reason is that job-seekers may have *distributional preferences*, i.e. they care about both their own income and their income relative to others. There is extensive evidence that interpersonal comparisons can influence individuals' sense of well-being (Frank, 1985; Loewenstein et al., 1989). Indeed, a large literature has emerged studying the impact of relative wage concerns on labour market outcomes (Akerlof and Yellen, 1990; Frank, 1984). As a consequence, economists have developed formal models incorporating distributional concerns into individual preferences. Using a prominent example of such a model (Fehr and Schmidt, 1999) and applying it to a job-search context, we show that the reservation wages of job-seekers with distributional preferences will be influenced by the wages other job-seekers have accepted.

A second reason that social comparisons may affect job search stems from *learning*. There is a large literature empirically supporting the idea that people learn from each other in complex environments (see Manski, 2000). Even if offer processes are independent,

[1]There is a large empirical literature about how individuals use networks to collect information in order to find jobs and how search outcomes are influenced by social connections. According to an overview by Ioannides and Loury (2004), the literature finds that the use of networks in job search is widespread and generally productive, as it increases the job-finding rate and improves the quality of the match. These findings are reinforced by a number of more recent studies (e.g. Brown et al., 2016; Cappellari and Tatsiramos, 2015; Cingano and Rosolia, 2012; Dustmann et al., 2016). A number of studies have developed theoretical frameworks for analysing the effects of social networks on job search (e.g. Calvó-Armengol and Jackson, 2004; Galenianos, 2013; Horváth, 2014). The current study suggests that social comparison effects offer an additional (and overlooked) explanation for the empirical phenomenon that search outcomes are correlated within social networks.

and even if job-seekers care only about their own narrowly-defined self-interest, in an environment where weighing up the attractiveness of a wage offer requires comparing it with a dynamic lottery, a job-seeker may be uncertain about whether it is worth accepting the wage offer. Whether others have accepted or rejected similar offers may then influence their decision. For example, if job-seekers observe others rejecting offers, this may lead them to reject similar offers.[2]

While field data offers ample evidence that social context matters because agents' behaviour within groups is correlated, it is notoriously difficult to identify whether this arises because of "peer effects", whereby individuals are affected by one another's behaviour (Manski, 1993). In addition, field data does not typically allow separate identification of network versus social comparison effects. For example, Marmaros and Sacerdote (2002) examine the link between social interaction and job search among Dartmouth college graduates. They find that employment outcomes for a student are correlated with mean outcomes for that student's freshmen hallmates and dormmates. Because hallmates and dormmates are randomly assigned, these correlations cannot be attributed to selection effects. These correlations may be due to some sort of peer effect, but (as noted by Marmaros and Sacerdote) they may also arise because of shocks which are correlated within groups. Moreover, even if the correlations reflect peer effects there is insufficient information to know whether these are driven by network or social comparison effects. The difficulties of disentangling network and social comparison effects using naturally occurring data motivated us to use an experimental approach.

Our framework for examining whether social comparisons affect job search is a sequential search decision experiment. This is based on McCall's (1970) "Basic Search Paradigm" in which searchers receive a wage-offer from a distribution and decide

[2]Distributional preferences and learning do not exhaust the possible mechanisms by which social comparisons affect behaviour. A large literature in psychology on "social facilitation" (Zajonc, 1965) emphasise the importance of being observed for performance on a variety of tasks.

whether to accept it, receiving this wage in each remaining period, or reject it and wait for another offer. We vary across treatments whether subjects make decisions in isolation or whether they can observe information from a partner's wage offers, decisions, and earnings. A long-established experimental literature has studied the sequential search task and its variants. Early studies compared individual behaviour with optimal search rules and alternative heuristics (Braunstein and Schotter, 1982; Hey, 1982, 1987; Oaxaca and Cox, 1989, 1992; Schotter and Braunstein, 1981) and more recent work has focused on behavioural explanations for deviations from theoretical predictions (Brown et al., 2011; Kogut, 1990; Schunk, 2009; Schunk and Winter, 2009; Sonnemans, 1998). A feature of all of this literature is that the search decisions are made in isolation. We extend this literature by adding a social context to the search task.

Our research question is whether social comparisons have an impact on job search, and so our experiment eliminates the possibility of network effects by design. Within the lab, we are able to tightly control the offer process, the interaction between subjects, and the information available to subjects. This allows us to exclude social network effects — subjects have no way of influencing another's offer process — and also to exclude common shocks to the offer process. To test for social comparison effects we vary across treatments the feedback given to subjects about the search decisions and outcomes of others. In a control *Isolated treatment* subjects conduct the search task in isolation. They receive no information about others' search decisions or outcomes. In two other treatments, subjects are randomly and anonymously paired and receive feedback on their partner's search decisions and outcomes as they conduct the search task. In the *Partner treatment* subjects observe their partner's income, whereas in the *Partner-Offer treatment*, they are also informed of the offers rejected by their partner. In all treatments the subjects are given full information about the stochastic offer process, and the earnings of a subject depend only on how they react to their own offers, and are independent of

the decisions and earnings of other subjects. Thus, for a fully rational self-interested subject, feedback about a partner's search decisions and outcomes is irrelevant.

In our Isolated treatment, in line with the long-established literature, average reservation wages are lower than the expected earnings maximizing benchmark, and reservation wages decline within the search task. In the Partner and Partner-Offer treatments we observe similar average levels and rates of decline of reservation wages, but we also observe that subjects systematically respond to feedback about their partner's search outcomes.

In the Partner and Partner-Offer treatments, when a subject observes their partner accepting a low wage they reduce their reservation wage more than they otherwise would, while if they observe their partner accepting a high wage they reduce their reservation wage less than they otherwise would. These results are qualitatively consistent with predictions of a model in which job-seekers care about earnings differentials as well as their own earnings. The Partner-Offer treatment also allows us to see whether the offers rejected by the partner affect reservation wages. These offers do not directly affect partner income, but do afford information about the partner's reservation wage. In fact, rejected partner offers do not influence the reservation wage, which suggests the within-task social comparison effects do not arise because job-seekers simply attempt to copy their partner's reservation wage. We do however see evidence of learning across search tasks. Although in all treatments initial reservation wages increase with experience, the increase is significantly greater in the treatments that allow social comparisons. We find that, in the social comparison treatments, subjects who are paired with a more successful and more patient partner in the previous task tend to increase their initial reservation wage in the next task.

In summary, our experiment reproduces the qualitative findings of previous studies that have focused on individual decision-making in isolated search environments. In

addition, in a first experiment studying this search task in a social context, we find systematic social comparison effects. Some of these effects are consistent with those predicted by a simple model of distributional preferences, and in addition we find that social feedback influences learning across repeated tasks.

The remainder of the paper is organised as follows. In Section 2.2 we review the (separate) relevant literatures on search experiments and social comparisons. In Section 2.3 we describe our experimental design and procedure, and in Section 2.4 we present our results. Section 2.5 offers a discussion and some concluding comments.

2.2 Literature

2.2.1 Search Experiments

The sequential search model (McCall, 1970) provides a tractable framework to analyse search decisions. The model assumes that a job-seeker receives one wage offer per period of search, where offers are independently drawn from a stationary and non-degenerate distribution, with a fixed outside option (typically a per-period unemployment benefit minus search cost). If the job-seeker accepts an offer they exit search and receive this wage for all future periods. If they reject the offer they wait for another offer in the next period. It is assumed that the time horizon is infinite and the job-seeker maximises expected discounted utility. The optimal decision is to accept any offer above or equal to a critical value (the reservation wage), which is constant over periods, and to continue with search otherwise.

Early experimental studies by Schotter and Braunstein (1982; 1981) and Hey (1987), examine how subjects search in a lab environment based on this basic search paradigm. Other studies have investigated variants of the model, such as using a fixed time horizon (Oaxaca and Cox, 1989, 1992), imperfect information about the offer distribution (Cox

and Oaxaca, 2000; Hey, 1981, 1982; Moon and Martin, 1990, 1992, 1996), heterogeneous search ability (Falk et al., 2009*a*,*b*) and variable search costs (Harrison and Morgan, 1990; McGee and McGee, 2016).

Previous experiments consistently show that subjects tend to (i) search too little and accept lower offers compared to the risk-neutral prediction, and (ii) reduce their reservation wages over the course of a sequential search task. Some behavioural theories have been proposed as explanations for these findings and have been tested in the laboratory. Studies suggest that the sunk-cost fallacy (Kogut, 1990; Sonnemans, 1998, 2000), loss and risk aversion preferences combined with bounded rationality (Schunk, 2009; Schunk and Winter, 2009; Soetevent and Bruzikas, 2016) and subjective costs of uncertain waiting time (Brown et al., 2011) all contribute to these regular departures from the theoretical predictions.

Some experiments repeat the search task, allowing investigation of how behaviour changes with experience, yet findings about learning effects are mixed. Kogut (1990) is the first experiment which repeats the search task within subject to allow for experience effects. With a small sample of subjects who conduct a series of repeated search tasks, he identifies a significant learning effect (i.e. search duration increases when subjects become more experienced) for half of the subjects. Sonnemans (1998) offers a more systematic examination of learning, showing that among subjects who repeat the search task many times, experimentation with different search strategies decreases with experience and search efficiency increases with experience. With fewer repetitions, Brown et al. (2011) find no significant learning effects.

All of these studies focus on individual decision-making, and exclude any social interaction. As far as we are aware the only experimental study embedding the basic search task in a social context is Ibanez et al. (2009), although their focus is rather different from ours. They examine whether search outcomes are improved if decisions

are made by a group instead of an individual, finding that decisions made by groups are indistinguishable from individual decisions. Our focus instead is on whether an individual's search decision is influenced by social comparisons.

2.2.2 Social Comparison Effects

We define social comparison effects to be the causal effect on an individual's decision from observing other individuals' independent decisions and outcomes. In the context of job search, independence means that the job-seeker's offer process and earnings are not directly affected by other job-seekers' decisions and outcomes. We are particularly interested in two mechanisms for social comparisons which are relevant in the context of job search: distributional preferences and learning.

Field and experimental studies have found social comparison effect in various contexts, such as labour effort provision (Abeler et al., 2010; Angelova et al., 2012; Clark et al., 2010; Cohn et al., 2014; Gächter et al., 2012, 2013),[3] job satisfaction (Card et al., 2012; Clark and Oswald, 1996), ultimatum and dictator games (Bicchieri and Xiao, 2009; Bohnet and Zeckhauser, 2004; Knez and Camerer, 1995; Krupka and Weber, 2009) and choice under risk (Friedl et al., 2014; Gamba et al., 2017; Haisley et al., 2008; Lahno and Serra-Garcia, 2015; Linde and Sonnemans, 2012, 2015; Rohde and Rohde, 2011; Schwerter, 2016).

Distributional preferences have been frequently proposed as an explanation for social comparison effects. As we show later, applying Fehr and Schmidt's (1999) model of distributional preferences to our laboratory setting leads to specific reference wage effects. However, there is mixed evidence on the force of distributional preferences when choices do not directly affect the earnings of others. Rohde and Rohde (2011), Linde and Sonnemans (2012), Lahno and Serra-Garcia (2015) and Schwerter (2016) find that

[3] By contrast, Charness and Kuhn (2007) find that effort choices are not affected by co-worker's wage (although the majority of subjects claim that they care about the co-worker's wage).

individual choices are sensitive to the earnings of other subjects even when subjects have no direct effect on each others' earnings. Linde and Sonnemans (2015), however, find no difference between individual decisions made in isolation and those made in a social setting.

Another potential source of social comparison effects is social learning. Subjects may be uncertain about the quality of their decisions, and therefore "may seek to draw lessons from observation of the actions chosen and outcomes experienced by others" (Manski, 2000). This might manifest itself in imitation: evidence for this is provided in a number of field experiments (Cai et al., 2009, 2015; Conley and Udry, 2010). Imitation might be a reasonable rule of behaviour if a subject believes others have better strategies. In a field experiment among investors, Bursztyn et al. (2014) find that investors are influenced by their peers, and financially unsophisticated investors tend to be affected more than the sophisticated ones. The extent to which people imitate might also depend on the observed performance of their peers. A particularly plausible rule is that people imitate those who perform better: Schlag (2011) refer to this as an "Imitate if Better" strategy. For instance, Offerman and Schotter (2009) show that when offered opportunities to sample the choices of others, subjects do choose to sample, and imitate the behaviour which leads to higher realised profits.[4]

Our paper is the first to investigate the relevance of distributional preferences and learning in a search task in which social comparisons are possible.

[4]The effects of getting additional search information about job market conditions have been investigated in field experiments by Altmann et al. (2016) and Belot et al. (2016). This is different from what we mean by social learning: the information in their settings comes from a centralised source (e.g. employment service) rather than from social interactions.

2.3 Experimental Design

2.3.1 The basic search task

We start by describing the search task which is common to all three treatments. To mimic the stationary, infinite horizon structure of the basic search paradigm, we use a random termination method whereby the search task terminates after each period with a constant probability.[5] The continuation probability in the uncertain horizon model we use plays the same role as the discount factor in the infinite horizon model: a subject maximizing expected earnings should use the same (constant) reservation wage as an infinitely-lived risk-neutral agent with a discount factor equal to the continuation probability.

The search task consists of a sequence of periods. In the first period, subjects choose a reservation wage. This is an integer between 0 and 1000. Subjects then receive an offer, randomly drawn from a discrete uniform distribution over 1 to 1000. The offer is automatically accepted if it is greater than or equal to the reservation wage, and rejected otherwise.[6] If the offer is accepted, subjects receive that wage as (points) income in all subsequent periods, receive no further offers and make no further decisions for the remainder of that sequence. If the offer is rejected, their income from the current period is zero. At the end of the period, subjects receive feedback which varies across the treatments, as we will discuss in the next subsection. The sequence continues with probability 0.95. If the sequence continues, subjects who have not accepted an offer are required to enter another reservation wage and the process repeats. At the end of the

[5] Random termination has been used extensively to mimic infinite-horizon games. See, for example, Dal Bó and Fréchette (2017). The only other search experiment of which we are aware which uses random termination is Brown et al. (2011).

[6] Subjects are required to enter the reservation wage within 15 seconds. If they do not enter a reservation wage within this time, this period's offer is rejected. The time restriction forces subjects to complete a period at roughly the same time, facilitating the synchronization of feedback. In the experimental instructions we refer to the reservation wage as a "minimum acceptable offer"; see instructions in Appendix 2.A for details. We have subjects enter reservation wages rather than make binary accept/reject decisions in order to obtain more information about their search strategy. This technique was first applied in search tasks by Oaxaca and Cox (1992).

sequence, subjects are informed of their points earnings from the sequence, which is simply the sum of points earnings from all periods in the sequence. The information received at the end of each sequence does not vary across treatments.

Subjects complete 10 sequences, each sequence having the same structure described above. All offers are independent draws across subjects, periods and sequences. At the end of the tenth sequence one sequence is randomly chosen for each session and points earnings from this sequence are used to determine subjects' monetary earnings. Subjects receive a show-up fee of 5.00 GBP plus additional earnings from the chosen sequence, converted at a rate of 3,000 points = 1.00 GBP.

All the information above is described carefully to subjects in the instructions which were handed out to subjects and read out at the beginning of the experiment. After reading out the instructions, control questions were given to ensure that subjects understood essential information about the search environment. See Appendix 2.A for a copy of instructions and control questions.

2.3.2 Treatments

To allow us to study the effect of social comparison, we manipulate across treatments the feedback received by subjects at the end of each period. In the simplest Isolated treatment, the feedback received by subjects in each period consists only of information on their own offer, whether it is accepted, the resulting per-period income and their accumulated earnings across all periods in that sequence. In the Partner and Partner-Offer treatments subjects are randomly re-paired at the beginning of each sequence. The random termination of search sequences, which is implemented independently for each subject in the Isolated treatment, is implemented independently at the pair level in the Partner and Partner-Offer treatments. In all three treatments, offers are independent across subjects and an additional control question is added in the Partner and Partner-

Offer treatments to ensure that subjects understand that own and partner offers are drawn independently.

In the Partner treatment subjects are additionally shown their partner's income at the end of each period. They are therefore aware of whether or not their partner has accepted an offer, and the amount of that offer if it was accepted. The Partner treatment allows for relatively little social feedback because subjects only learn their partner's income each period. Subjects whose partner has an income of zero cannot tell whether this is because their partner has a high reservation wage, or because their partner received low offers. This also means that there is relatively little within-sequence variation in the feedback received. The partner's per-period income is either zero (meaning the partner has not yet accepted an offer) or it is constant across periods.

In the Partner-Offer treatment, subjects again receive information about partner income at the end of a period. In addition, subjects are also shown the offer made to their partner, and whether or not their partner accepted or rejected that offer. Thus the Partner-Offer treatment allows for more feedback and more within-sequence variation. It also allows us to test whether feedback on rejected offers, as well as income, affects search behaviour.[7]

The main features of the design are summarised in Table 2.1. Each session consisted of 12 subjects. We conducted six sessions for each treatment. No subject participated in more than one session.

2.3.3 Discussion of the design

The Isolated treatment provides a benchmark which we can compare to the existing results from experimental analyses of search, all of which study decisions made in isolation. In this setting, expected points earnings are maximised by choosing a constant

[7] See Figure 2.A1 in Appendix 2.A for an illustration of the feedback screens seen by subjects in our three treatments.

Table 2.1 Treatment Design

	Isolated	Partner	Partner-Offer
Sessions	6	6	6
Subjects per session	12	12	12
Sequences per subject	10	10	10
Paired subjects	No	Yes	Yes
Mutual feedback on			
Offer	No	No	Yes
Income	No	Yes	Yes

reservation wage $r^* = 725$ (see Appendix 2.B.1). This is also the optimal reservation wage in the Partner and Partner-Offer treatments if subjects care only about expected earnings.[8] If social comparisons matter, the optimal reservation wage in the Partner and Partner-Offer treatments may differ from the Isolated treatment. Our experimental design allows us to test the two different potential social comparison mechanisms discussed in the introduction: distributional preferences and learning.

If subjects have distributional preferences this will affect the optimal reservation wage. Specifically, following Fehr and Schmidt (1999), we assume that subject i dislikes income differences, as represented by the period utility function

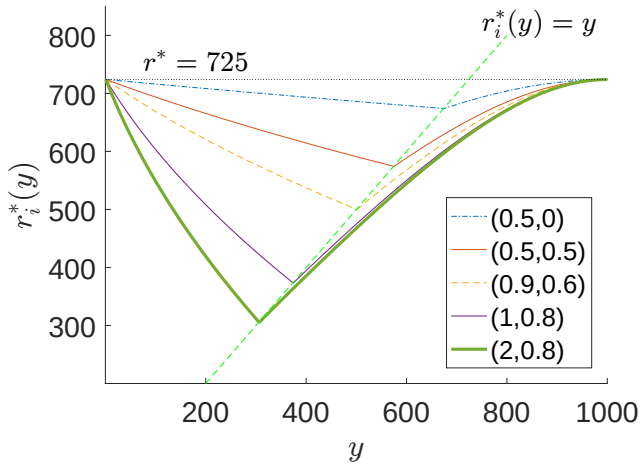
$$u_{it}(y_{it}, y_{jt}) = y_{it} - \alpha\big(\max\{y_{it}, y_{jt}\} - y_{it}\big) - \beta\big(\max\{y_{it}, y_{jt}\} - y_{jt}\big), \tag{2.1}$$

where y_{it} denotes i's period t income, y_{jt} denotes the period t income of i's partner, α is the marginal disutility from disadvantageous inequality and β is the marginal disutility from advantageous inequality.

[8] If a subject is risk averse, the level of the optimal reservation wage will change, but will still be constant across periods, and will still be the optimal reservation wage regardless of which treatment the subject is in.

If i's partner has exited with an income of $y_{jt} = y$, the optimal reservation wage $r_i^*(y)$ can be shown to be a decreasing function of the partner's income if the partner has accepted a low wage, and an increasing function if the partner has accepted a high wage (see Appendix 2.B.2 for details). We plot the optimal reservation wage function in Figure 2.1 for five sets of (α, β), chosen based on the assumptions made in Fehr and Schmidt (1999), i.e. $0 \leq \beta \leq 1$ and $\alpha \geq \beta$.

Fig. 2.1 Optimal reservation wage for i after j exits with y: $r_i^*(y)$



Notes: (α, β) for each curve is given by the legend; the horizontal dotted line indicates the expected earnings maximizing $r^* = 725$; the kink point of $r_i^*(y)$ lies on the 45-degree line.

To identify the relevant portion of the reservation wage function, we need to analyse the dynamic game where both i and j are still searching. Identifying the equilibrium initial reservation wage requires assumptions about both subjects' preferences and what they know about each other's preferences. For tractability, we assume subjects have identical preferences and this is common knowledge. In Appendix 2.B.2 we show that the initial equilibrium reservation wage lies above the kink point of the $r_i^*(y)$ function. This implies that a subject never accepts a wage below the kink point and so in equilibrium only the positively-sloped part of the reservation wage function is relevant in Figure 2.1.

The initial reservation wage is very close to the expected earnings maximizing reservation wage r^* and remains constant until one of the subjects accepts an offer (i.e. rejected offers have no effect on reservation wages). If the partner exits, the subject adopts a new, constant reservation wage. The new reservation wage is increasing in the partner's accepted offer, but in expectation is lower than the initial reservation wage. We refer to the positive relationship between a partner's income and the subject's reservation wage as a reference wage effect.

To test for this reference wage effect in our experimental data we examine the change in the reservation wage within sequences. We examine how the reservation wage changes following the partner's exit, and whether this change is moderated by the partner's income. The change in the reservation wage is predicted to be the same in both the Partner and Partner-Offer treatments, while the Isolated treatment provides a natural placebo test.

Within-sequence social comparison effects can also arise from a desire to imitate the partner's decision. In the Partner and Partner-Offer treatments, the feedback allows the subjects to make (limited) inferences about their partner's reservation wage.[9] If they observe the partner accept an offer they would have rejected, they can infer that their partner had a lower reservation wage. If the subject only desires to imitate their partner, they should reduce their reservation wage, and their new reservation wage will be no greater than the partner's accepted offer. In the Partner-Offer treatment there is an additional scenario that would induce an imitative subject to change their reservation wage. If the subject observes the partner reject an offer they would have accepted, they can infer that their partner had a higher reservation wage. Thus they should increase their

[9]We do not directly reveal subjects' reservation wages to their partners. This reflects the fact that, in real job search environments, job-seekers do not typically observe their peers' reservation wages, but they often do observe their peers' search outcomes (as in the Partner treatment) or the offers their peers receive (as in the Partner-Offer treatment).

reservation wage, and their new reservation wage will be no smaller than the rejected offer.

The reference wage effect and the imitation effect are largely observationally equivalent, and consequently difficult to disentangle. However, with our design, the Partner-Offer treatment offers a testable difference: the rejected offers of partners have no impact according to the reference wage effect, whereas they have a positive impact according to the imitation effect.

The fact that we have repeated search sequences also allows us to test for an effect of social feedback across sequences. In the Isolated treatment, subjects can learn about how their own strategies affect sequence earnings and, based on prior literature (see Section 2.2.1), experience from earlier sequences may improve search efficiency in later sequences. In the social feedback treatments, the additional information about the partner's strategies and income might be expected to enhance this process. To examine this, we test whether behaviour is consistent with a simple learning strategy, whereby subjects imitate their partners from the previous sequence if their partners had higher income. This strategy implies that subjects whose partners have higher income and longer search durations will increase their reservation wage by a greater amount.

2.4 Results

216 subjects who were recruited through ORSEE (Greiner, 2015) participated in the 18 sessions. None of the subjects had participated in a search experiment before. All sessions lasted for approximately 50 minutes, and earnings averaged 10.11GBP per subject inclusive of the show-up fee. The experiment was computerised using the z-Tree program (Fischbacher, 2007) and conducted in the CeDEx laboratory at the University of Nottingham in March and April 2015 and March 2016. Subjects were randomly

assigned into Isolated, Partner and Partner-Offer treatments. Observable characteristics are similar across the treatments, as shown in Table 2.2. The differences in proportions are insignificant according to pooled chi-square tests (p-values are reported in the last column).

We observe 2,160 sequences of search (3 treatments $\times$ 6 sessions $\times$ 12 subjects $\times$ 10 sequences), which contain 4,749 decision-making periods.[10] 5.79% of sequences were censored, which means the subject was still searching when the sequence was terminated by the computer and no offer was accepted in that sequence.

Table 2.2 Observable characteristics of subjects across treatments

Treatment	Isolated		Partner		Partner-Offer		p-value from χ^2 test
Female	45	(62.5%)	44	(61.11%)	49	(68.06%)	0.656
Postgraduate	6	(8.33%)	6	(8.33%)	6	(8.33%)	1.000
Field of study							
Science	34	(47.22%)	37	(51.39%)	39	(54.17%)	0.703
Social Science	30	(41.67%)	22	(30.56%)	22	(30.56%)	0.268
Arts	8	(11.11%)	13	(18.06%)	11	(15.28%)	0.498

Notes: Sciences include Engineering and Medicine and Health Sciences. Double majors are classified as Sciences if a science major is combined with another major, or as Social Sciences if a social science major is combined with arts major.

Table 2.3 summarises levels and changes in reservation wages for each treatment. It shows that the two regularly-observed departures from the theoretical predictions discussed in Section 2.2.1 are replicated in all treatment conditions of our experiment. First, in all treatments average reservation wages are below the reservation wage of an expected earnings maximiser ($r^* = 725$). Panel (a) of Table 2.3 reports average initial (period 1) reservation wages of 539, 555 and 512 in the Isolated, Partner and Partner-Offer treatments respectively, with 95% confidence intervals easily excluding r^*.[11] As a result of these low reservation wages, search sequences tend to be short (the

[10] In 7 periods, the subject passed the 15-second time limit and consequently the offer was rejected. This happened to 6 different subjects, 1 in the Isolated, 1 in the Partner and 4 in the Partner-Offer treatment.

[11] Throughout the paper, standard errors are clustered at the session level.

Table 2.3 Overview of reservation wages across treatments, pooled across sequences

	Obs.	Mean	Std. Err.	95% C.I.	
(a) Initial period					
Isolated	719	538.59	23.08	(479.26	597.91)
Partner	719	555.07	40.36	(451.33	658.81)
Partner-Offer	717	512.00	10.61	(484.71	539.28)
(b) Overall average					
Isolated	1,575	586.85	14.51	(549.56	624.15)
Partner	1,587	585.42	29.03	(510.80	660.04)
Partner-Offer	1,580	543.10	15.01	(504.53	581.68)
(c) Within sequence per-period change					
Isolated	855	−26.50	4.53	(−38.15	−14.85)
Partner	867	−29.81	3.65	(−39.18	−20.43)
Partner-Offer	859	−32.47	3.24	(−40.81	−24.13)

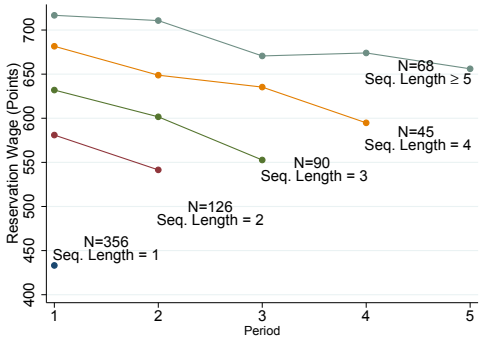
Notes: 7 periods in which the subject fails to enter a decision before the time limit (0.1% of active periods) are excluded from the summary.

average duration of search in our treatments is 2.19 periods), and subjects often exit with an offer below r^* (37.45% of the sequences in our sample). Average reservation wages in all periods summarised in Panel (b) are slightly higher than average initial reservation wages, mainly due to a selection effect (subjects who set lower reservation wages tend to exit search earlier and so are under-represented in the sample), but still significantly below r^*. This finding is consistent with the existing experimental literature. For example, our average reservation wage is just below 80% of r^*, and a similar ratio can be calculated from the average reported in Brown et al. (2011).

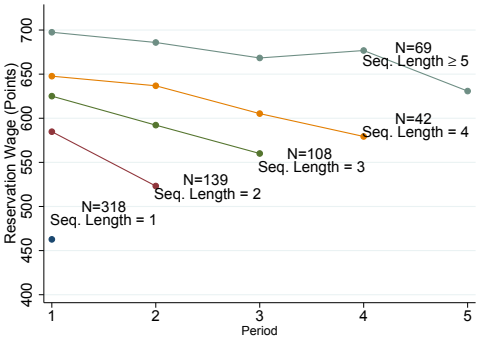
Second, in all treatments we consistently observe a declining reservation wage within sequences. Panel (c) of Table 2.3 shows that the average per-period change in the reservation wage is -27, -30 and -32 in the Isolated, Partner and Partner-Offer treatments. The 95% confidence intervals all exclude zero. Figure 2.2 shows this graphically by plotting the average reservation wage across periods, conditional on search length.[12] For all subsamples, the reservation wage is declining in all treatments.

[12]We condition on search length because subjects with higher reservation wages will tend to have longer search durations. If we did not condition on search length, the average reservation wage would

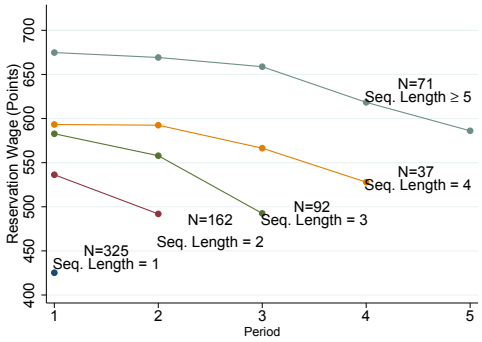
Fig. 2.2 Reservation wage by search length



(a) Isolated treatment



(b) Partner treatment



(c) Partner-Offer treatment

Notes: *N* in each figure refers to the number of sequences with that search length. To ensure that sequences are not selected into the subsamples by random termination, 106 sequences (4.91%) which are censored due to termination before the fifth period are not included in the summary. The sequences in which the subject passed the time limit for the initial period are also excluded.

This finding is also consistent with the existing experimental literature, see for example Braunstein and Schotter (1982) and Brown et al. (2011).[13]

At an aggregate level, we find strong similarities across treatments. Neither initial period reservation wages, average reservation wages across all periods, or the within

increase across periods as subjects with lower reservation wages exit faster, even if subjects have constant reservation wages.

[13]Indirect evidence of the declining reservation wage also comes from studies which do not elicit a reservation wage but allow subjects to recall offers. In these studies, many subjects exercise the option to recall offers that have been previously rejected (Hey, 1987; Kogut, 1990).

sequence per-period changes in reservation wages are significantly different across treatments.[14]

However, in the next two sections we analyse the data at a more disaggregated level and show that subjects systematically respond to social feedback in our Partner and Partner-Offer treatments, both within and between sequences.

2.4.1 The impact of social feedback within sequences

To test whether subjects are responding to their partner's search feedback, we compare the decisions made before and after feedback is received. To do this, we focus on the difference between the reservation wage in period t and the reservation wage in the previous period $t-1$. Search feedback from the partner is characterised by two key variables. First, an indicator of whether subject i's partner j accepted an offer in the previous period. Second, the value of the offer received by j in the previous period. These variables are only defined in periods in which subject i is still searching at t, and their partner j was still searching at $t-1$, and therefore the sample used in the analysis in this section consists of those observations from periods t in which the subject's partner was still searching at $t-1$. In the Isolated treatment subjects do not have a partner, we therefore create artificial partners by pairing subjects randomly. The artificial partners do not provide any feedback, and hence are used for placebo tests.

There are 1,678 periods in the sample, which is 35.19% of all the decision-making periods. The initial periods of a sequence (2,160, 45.48%) are excluded because no feedback from the partner is possible. Similarly, periods in which the partner exited before $t-1$ (910, 19.16%) are excluded because, again, no feedback is possible at t.[15]

[14] Treating each subject as an independent observation ($N = 3 \times 6 \times 12$), a Kruskal-Wallis rank test gives $p = 0.177$ for equality of initial period reservation wages, $p = 0.139$ for equality of overall average reservation wages and $p = 0.764$ for equality of per-period change in reservation wages.

[15] In addition, 8 periods (0.17%) in which the subject timed out in the current or the previous period are also excluded.

Almost all subjects (96.76%) are included in the sample for one or more sequences (the average is 5.28 sequences per subject), and observations are fairly evenly distributed across sequences.

In Figure 2.3 we plot the correlation between the change in subject i's reservation wage and the value of the offer received by their partner. We do this separately, by treatment, for those cases where j rejected the offer (top row) and those cases where j accepted the offer (bottom row). We separate in this way because accepted offers lead to a change in partner's income, whereas rejected offers do not. The sample size, the Pearson correlation coefficient and p-value are given at the top of each panel, and in addition we display a scatter plot of the average change in the reservation wage and average value of the partner's offer for each decile of the partner's offer (10 data points) and a linear fit. As noted earlier, Panels (a) and (d) are placebo conditions from the Isolated treatment. Panel (b) is also a placebo condition because there is no feedback on rejected offers in the Partner treatment. As expected, correlations in all the placebo conditions are insignificant. The correlation between the change in reservation wage and partner's rejected offer is also small and insignificant in the Partner-Offer treatment (Panel (c)). The patterns are quite distinct in those cases in the Partner and Partner-Offer treatments where the partner accepts the offer (Panels (e) and (f)). The change in reservation wage is significantly and positively correlated with partner's offer, with a correlation coefficient of 0.144 ($p = 0.025$) and 0.179 ($p = 0.006$).

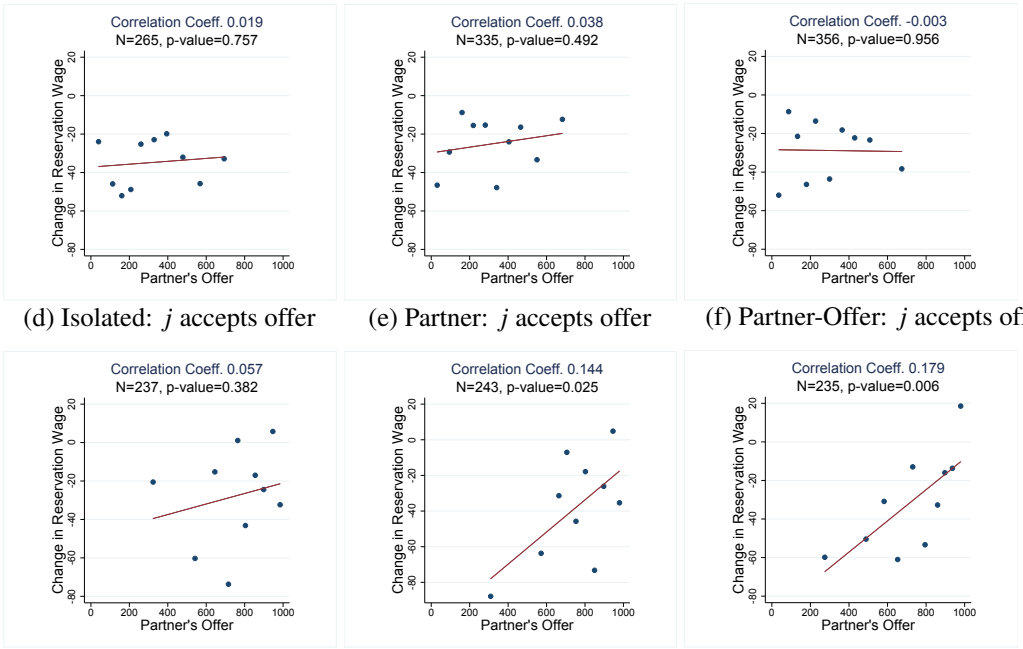
To quantitatively estimate how subjects are influenced by social comparisons, we estimate the following model for each treatment:

$$r_{is,t} - r_{is,t-1} = \beta_0 + \beta_1 a^j_{is,t-1} + \beta_2 (a^j_{is,t-1} \cdot w^j_{is,t-1}) + \beta_3 w^j_{is,t-1} + \lambda_t + \theta_s + \eta_i + \varepsilon_{is,t} \quad (2.2)$$

The subscripts i, s and t are identifiers for, respectively, subject, sequence and period. The dependent variable is the change in the reservation wage relative to the previous

Fig. 2.3 Correlation between changes in reservation wage and partner's offer

(a) Isolated: j rejects offer (b) Partner: j rejects offer (c) Partner-Offer: j rejects offer



(d) Isolated: j accepts offer (e) Partner: j accepts offer (f) Partner-Offer: j accepts offer

Notes: We separate the observations in each treatment by the decision of the partner (or placebo partner for the Isolated treatment). The Pearson correlation between partner's offer and changes in reservation wage, and unadjusted, significance level are given at the top of each panel. The panels display the binned scatter plots: for each decile of partner's offers in the subsample, we plot the average change in the reservation wage and the average of partner's offer. Also shown is the fitted linear relationship among the 10 dots. Details of this method can be found in Stepner (2013).

period. The explanatory variables are an indicator for whether the partner accepted their offer in period $t-1$ ($a^j = 1$ for acceptance and 0 for rejection), the value of the partner's offer in period $t-1$ (w^j) and the interaction of the acceptance indicator and the offer, which is the period $t-1$ income of the partner. We control for period (λ_t) and sequence (θ_s) fixed effects, and allow for subject random effects (η_i).

In the Isolated treatment a^j and w^j are not observed, so changes in the reservation wage must be independent of these variables by design, and so β_1, β_2 and β_3 should be zero. In the Partner treatment a^j is observed but w^j is only observed if $a^j = 1$, so β_3 should be zero. In the Partner-Offer treatment, both a^j and w^j are observed in every

period and so β_1, β_2 and β_3 might all be different from zero if subjects respond to this information.

Estimates of (2.2) are reported in Table 2.4.[16] On average across all periods and sequences in the sample, subjects reduced their reservation wage by about 30 points per period. This is very similar to the per-period decline reported in Table 2.3 and Figure 2.2. In the Isolated treatment, as expected, $\hat{\beta}_1$, $\hat{\beta}_2$ and $\hat{\beta}_3$ are all small and insignificantly different from zero. In the Partner and Partner-Offer treatments, however, $\hat{\beta}_1$ is significant and negative, indicating that a subject whose partner exits in the previous period lowers their reservation wage by 60–70 points more than they otherwise would. $\hat{\beta}_2$ is significant and positive in the Partner and Partner-Offer treatments, showing that the value of the accepted offer is also relevant. A subject's reservation wage is increasing with the offer accepted by the partner; the slope is approximately 0.1 ($\hat{\beta}_2 + \hat{\beta}_3$). By comparison, in our analysis of job-seekers with Fehr-Schmidt preferences the slope varies between 0.13 and 0.34 depending on preference parameters (see Table 2.B1 in Appendix 2.B.2 for details).

The exit and income effects together imply that a subject whose partner exits may increase or reduce their reservation wage in the next period, depending on the value of the partner's exit wage. Given that the average exit wage of the partner is 746 (Partner treatment) and 718 (Partner-offer treatment), our estimates of model (2.2) predict a very small negative effect of partner exit on reservation wage. This leads to a small predicted increase in the probability of exit in the period after a partner exits of less than 1%.[17] By comparison, in our theoretical analysis the remaining job-seeker's average exit probability increases by 1% to 3% depending on preference parameters (also see Table 2.B1).

[16] The results are robust to alternative specifications, such as dropping the period and sequence dummies, or the addition of controls for observable characteristics, see Table 2.C1 in Appendix 2.C.

[17] This small effect helps to explain why the average change in the reservation wage is similar in the Isolated and social feedback treatments.

Table 2.4 Change in reservation wage between periods as a function of partner's offer and acceptance decision.

	(1) Isolated	(2) Partner	(3) Partner-Offer
Mean Δr	−31.538	−30.702	−29.844
β_1 j accepted offer (a^j)	−16.621	−71.205**	−63.645***
	(11.813)	(33.313)	(18.923)
β_2 j's income ($a^j w^j$)	0.029	0.092**	0.088***
	(0.026)	(0.046)	(0.027)
β_3 value of j's offer (w^j)	0.002	−0.006	−0.002
	(0.016)	(0.022)	(0.013)
Subject random effects	Yes	Yes	Yes
Period fixed effects	Yes	Yes	Yes
Sequence fixed effects	Yes	Yes	Yes
Number of obs.	502	578	591
Number of sessions	6	6	6

Notes: Standard errors are in parentheses.

In contrast to the estimated effect of accepted offers ($\hat{\beta}_2 + \hat{\beta}_3$), offers which are rejected appear to have no impact on subjects' reservation wages: $\hat{\beta}_3$ is small and insignificantly different from zero in all treatments.

These regression results reinforce the findings reported in Figure 2.3. First, all estimated coefficients in placebo conditions are insignificant. Second, the results reproduce the positive correlation between the partner's income and the change in the subject's reservation wage in the Partner and Partner-Offer treatments. Third, the results reproduce the insignificant effect of rejected offers on reservation wages in the Partner-Offer treatment. The latter two findings are qualitatively consistent with the Fehr-Schmidt model analysed in Appendix 2.B.2.

The second finding is also consistent with subjects inferring their partner's reservation wage and attempt to imitate it. As noted earlier, in our setting, subjects can make inferences about their partner's reservation wages when they observe their partner's income. When a subject observes their partner accept an offer below their own reservation

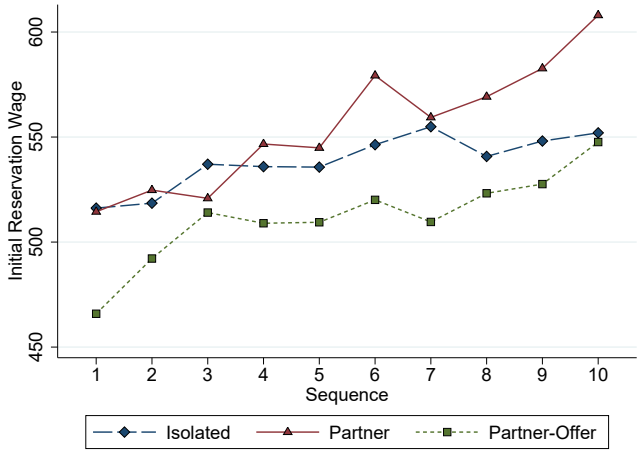
wage they can infer that the partner's reservation wage is lower than their own. If the subject wishes to imitate they will lower their own reservation wage. For a fixed initial own reservation wage the reduction will be larger the lower the accepted offer. However, reservation wage imitation would also imply an analogous rejected offer effect when the subject observes their partner reject an offer above their own reservation wage. We do not observe such an effect in either Figure 2.3c or Table 2.4.

2.4.2 Impact of social feedback between sequences

We now consider whether feedback also changes reservation wages between sequences. Recall that subjects and their partners are randomly re-paired at the end of each sequence, and only one randomly-drawn sequence is used to determine earnings. A sequence is therefore very much like previous sequences, and from this perspective we might expect behaviour to be similar across sequences. However, behaviour might change across sequences if subjects learn from their own outcomes and the outcomes of their partners in earlier sequences.

In Figure 2.4 we plot the average period 1 reservation wage in each sequence across all 10 search sequences. We use period 1 reservation wages for this comparison because reservation wages in period 1 are not yet influenced by the within-sequence feedback documented in the previous section. Figure 2.4 shows that reservation wages tend to increase across sequences in all three treatments, but this increase is larger for the Partner and Partner-Offer treatments. A linear trend estimated separately for each treatment yields a slope of 9 points (standard error 2 points, $p = 0.005$) per sequence in the Partner treatment, 6 $(3,\ p = 0.093)$ in the Partner-Offer treatment and 3 $(3,\ p = 0.227)$ in the Isolated treatment. This implies that, in the Partner treatment, subjects choose a reservation wage 16.63% higher in the final sequence compared to the first sequence. In the Partner-Offer treatment, this increase is estimated to be

Fig. 2.4 Period 1 reservation wages by treatment and sequence



11.71%. In the Isolated treatment, this increase is estimated to be 6.43%. However, the difference in sequence trend between Partner and Isolated treatments is only marginally significant ($p = 0.067$), and the difference between Partner-Offer and Isolated treatments is insignificant ($p = 0.499$).[18]

Figure 2.4 suggests that the additional feedback provided in the Partner and Partner-Offer treatments may cause subjects to increase their reservation wage more between sequences. We now investigate whether the extent to which reservation wages increase between sequences varies systematically with the outcomes of partners in previous sequences, in those treatments where feedback is available.

During each sequence, subjects in the Partner and Partner-Offer treatments observe their partner's per-period income and search duration. We consider an "Imitate if Better" rule, i.e. if the partner does better the subject move in the direction of partner's reservation wage, but if the partner does worse the subject simply carry on with the same reservation wage. To operationalise this, we define "better" to be a case where the partner achieved a

[18] The findings are robust to the inclusion of subject characteristics, reported in Table 2.C2. Table 2.C2 also examines sequence trends for exit period reservation wages. These are positive and significantly different from zero in all treatments, but are higher in the social feedback treatments (and again the difference is only significant between Partner and Isolated treatments).

higher per-period income.[19] Because subjects do not observe their partner's reservation wage directly, we assume that they infer it from their partner's search duration. We assume that subjects whose partner has a longer (shorter) search duration infer that their partner has a higher (lower) reservation wage. If this is the case, then a subject who observes a partner with a higher income and a longer (shorter) duration will tend to increase (decrease) their reservation wage between sequences.

To test this hypothesis, we measure the between-sequence change in period 1 reservation wages, $\Delta r_{is1} = r_{is1} - r_{i(s-1)1}$. Each subject-sequence is categorised in the following six types, according to the behaviour and outcome of their partner in the previous sequence:

Lower, Shorter j did not accept a better offer and search was shorter than i;

Lower, Same j did not accept a better offer and search was the same as i;

Lower, Longer j did not accept a better offer but search was longer than i;

Higher, Shorter j accepted a better offer and search was shorter than i;

Higher, Same j accepted a better offer and search was the same as i;

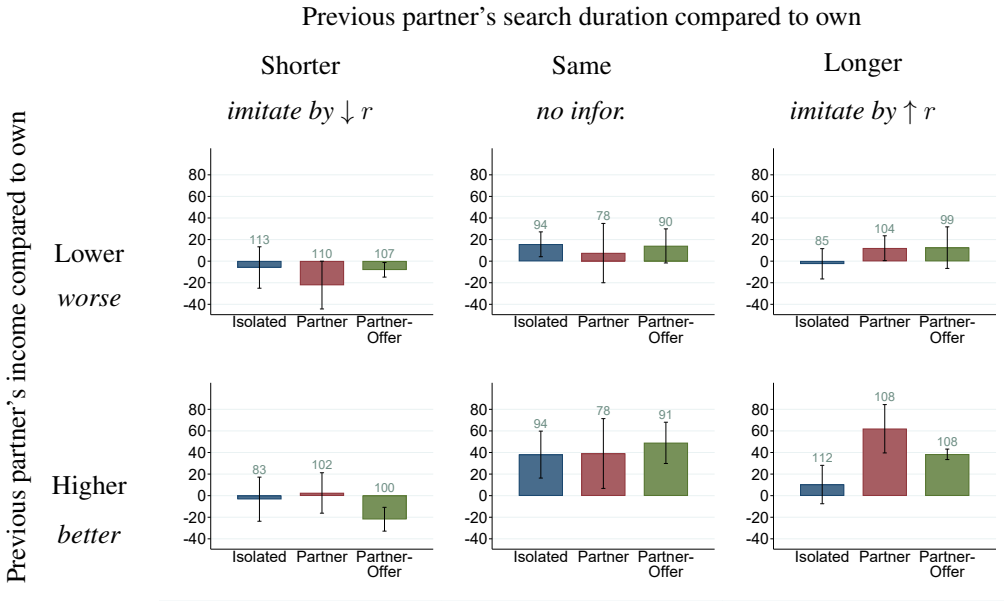
Higher, Longer j accepted a better offer and search was longer than i.

The selection of subjects into these categories is non-random even in the Isolated treatment with artificial Partners who provide no feedback. For example, subjects' reservation wages will, on average, be low if their partners have longer search duration. For this reason we compare Δr_{is1} between treatments within each of the six categories.

Figure 2.5 reports the average value of Δr_{is1} for each treatment and category. The top row of Figure 2.5 shows those categories where the partner had lower per-period income in the previous sequence than the subject. In none of these categories do we observe any

[19] We use per-period income rather than total earnings in the sequence because subjects receive direct feedback on partner's per-period income, but not total earnings in the sequence.

Fig. 2.5 Impact of partner's outcome in the previous sequence on the change in the reservation wage between sequences



Notes: Each bar indicates the mean value of Δr_{is1} in each treatment and category. The sample size for each cell is labelled at the top of the confidence interval. We exclude the 8.52% sequences for which the previous sequence is terminated before the subject or their partner accepts an offer.

significant difference in Δr_{is1} between treatments. Since in these cases the partner has done worse than the subject, this is all consistent with an "Imitate if Better" rule.

The bottom row of Figure 2.5 shows those categories where the partner had higher income than the subject in the previous sequence. The bottom centre panel shows the (Higher, Same) category. There is no significant difference in Δr_{is1} between treatments. The bottom right panel shows the (Higher, Longer) category. In this category, the "Imitate if Better" rule implies that the provision of feedback will cause subjects to increase their reservation wage between sequences. In the Partner treatment, subjects increased the reservation wage by 62 points between sequences when they were previously paired with a partner who accepted a higher offer and who searched longer in the previous sequence. This increase in the reservation wage is significantly higher than in the Isolated treatment

in the same panel ($p = 0.003$). The equivalent increase in the Partner-Offer treatment is 38 points ($p = 0.009$). The response in the Partner-Offer treatment is smaller than the Partner treatment (though not significantly so). This suggests that the additional feedback on Partner's rejected offers did not contribute to increased reservation wage in later sequences.

The bottom left panel shows the (Higher, Shorter) category. In this category, we predict that the provision of feedback will cause subjects to reduce their reservation wage between sequences. In the Partner treatment, we do not observe any reduction in the reservation wage, and there is no significant difference between the Partner and Isolated treatments ($p = 0.689$). In the Partner-Offer treatment there is a reduction in the reservation wage of 22 points, but again there is no significant difference between the Partner-Offer and Isolated treatments ($p = 0.135$).

In summary, our results provide mixed evidence for "Imitate if Better" behaviour. Subjects do appear to imitate partners with higher income who have longer search durations, but they do not imitate partners with higher income who have shorter search durations. Recall that, on average, subjects chose a reservation wage substantially below the expected-earnings maximizing benchmark of 725 (see Table 2.3). Thus, the way subjects respond to social feedback moves them closer to this benchmark in later sequences. This occurs because subjects whose partners searched longer and accepted higher offers were typically observing behaviour which was closer to the benchmark, and these subjects increased their own reservation wage by more in response.

2.5 Conclusions

It is widely recognised that social interaction matters for outcomes in labour markets. However, the existing literature has tended to focus on network effects, whereby social

networks provide additional options and information about those options. In this paper we use an experimental setting to investigate an alternative channel for social interactions to influence search: social comparison effects.

Our experiment employs a search task based on McCall's (1970) canonical model of labour market search, which offers predictions for optimal search decisions in a known search environment. Similar search tasks have been examined in the earlier experimental literature. To implement the stationary infinite-horizon decision-making problem, we use a random termination method whereby subjects search with an uncertain time horizon. In this setting, the continuation probability plays a similar role to the discount factor, and the strategy that maximises expected earnings is the same constant reservation wage that maximises the expected discounted utility in the infinite-horizon model.

In the treatment where subjects search in isolation, the findings from the previous literature are replicated in our search task with random termination. Subjects tend to set lower reservation wages than is consistent with maximizing expected earnings, and they tend to decrease reservation wages over the course of a search task.

We then consider how social comparisons affect search behaviour. We conduct two treatments in which subjects are paired and complete independent search tasks. During the task they receive information about their partner's search, but importantly their decisions have no direct effect on their partner's offer process. Thus our experiment excludes possible network effects.

In one treatment we allow subjects to observe a partner's income at the end of each period. This information is available both during search and after offer acceptance; these two states are analogous to the search and employment spells in a job-seeker's labour market experience. In another treatment we also allow subjects to observe offers received by a partner. This in addition to the partner's acceptance/rejection decision gives

additional information about the partner's reservation wage. Such social comparisons may affect search due to distributional preferences or learning effects.

In these novel treatments we show that social comparisons have a real impact on search decisions. First, we find that when a subject observes their partner exit with a low wage they tend to decrease their reservation wage by more than they otherwise would, and if they observe their partner exit with a high wage they tend to decrease their reservation wage by less than they otherwise would. Second, compared with our control treatment, we observe greater increases in initial reservation wages across search tasks in our treatments that allow social comparisons.

Some of these findings are qualitatively consistent with a reference wage effect driven by distributional preferences. For example, applying the Fehr-Schmidt model of inequality aversion to our setting, an active searcher's reservation wage increases with the income of a partner over a particular range of parameter values. However, social preference models cannot give a complete explanation for differences between our treatments. For example, they fail to account for stronger between-sequence changes in initial reservation wages in our treatments allowing social comparisons. This finding is better accounted for by a learning explanation whereby subjects tend to imitate the choices of more successful partners.

The social comparison effects we observe imply that job-seekers who search in a social environment will have an increased job-finding rate after their partner has found a job, because their reservation wage declines. Thus, the difference in search duration between partners will be smaller than in the case of two isolated job-seekers. The effects also imply that job-seekers who search in a social environment will have positively correlated wages even when the offer distributions they face are uncorrelated, because the partner's exit wage influences the remaining job-seeker's reservation wage.

An obvious question is whether the social comparison effects we observe are relevant to real job-search settings. We note that the social interaction in our experiment is very limited; the subjects in our various treatments experience a very similar environment, the only difference being limited to information about an anonymous other. In real job search environments we expect social interaction to be much richer. We see no reason to suppose that social comparison effects would be weaker in the richer environment. This question could be further explored in future research which varies the form of social interaction. We hypothesise that with richer interaction, for example allowing communication between subjects, we would observe stronger social comparison effects. We leave such issues to further research.

Appendix 2.A Instructions

[*] *indicates differences across treatments.*

Welcome! You are about to take part in a decision-making experiment. For participating in this experiment you will earn a show-up fee of £5. You can earn additional money depending on your decisions, so it is important that you read these instructions with care.

Please switch off your mobile phones. Communication with other participants is prohibited during the experiment. If you violate this rule, you will be dismissed from the experiment and you will forfeit all payments. If you have any questions, please raise your hand. A member of the experiment team will come and answer them in private.

The experiment will consist of 10 sequences. In each sequence you can earn points. After the completion of the experiment, the computer will randomly draw one sequence and all participants will receive additional earnings determined by the points they earned in that sequence. The points earned in that sequence will be converted to Pounds at the following rate:

3,000 points = 1 Pound

At the end of the experiment this amount plus your show-up fee of £5 will be paid to you privately in cash.

Description of a sequence

Each of the 10 sequences of the experiment has the same structure. Each sequence consists of multiple periods. Your total earnings from a sequence will be the sum of points you earn over the periods.

You will not know exactly how many periods there will be in a sequence until the sequence ends. The number of periods is determined by the computer following a simple rule: after each period, there is a 5% chance that the sequence ends and a 95% chance the sequence continues.

In period one, the computer will randomly draw an offer from 1 to 1000 points. Every offer is equally likely. Before you see the offer, you will be asked to choose a minimum offer that you would be willing to accept. The computer will then accept the offer for you if, and only if, it is above or equal to the minimum value you chose.

- If the offer is accepted, you earn that amount in period one and in every future period of the sequence.
- If the offer is rejected, your earnings from period one will be 0 and, if the sequence continues, you will receive a new offer in period two.

In period two, if you accepted the offer in period one you do not have to make a decision. Your earnings for the period will be equal to the offer you accepted in period

one. If you rejected the offer in period one, you will again be asked for the minimum offer you would accept, and you will receive a new offer, which again will be accepted if and only if it exceeds your minimum acceptable offer.

- If the offer is accepted you earn that amount in period two and in every future period of the sequence.
- If the offer is rejected, your earnings from period two will be 0 and, if the sequence continues, you will receive a new offer in period three.

This process will continue until the sequence ends. Remember there is a 95% chance that the sequence continues, and a 5% chance that it ends, at the end of any period. At the end of the sequence your earnings for the sequence will be the sum of your earnings from each period in the sequence.

As long as you have not accepted an offer in the sequence, if the sequence continues to the next period you will get a new offer. The new offer will be drawn independently from the previous offers. This means in every period the offers from 1 to 1000 are equally likely regardless of what offers were made previously.

In each period where you have to enter a minimum acceptable offer you are required to enter it within 15 seconds. If you do not enter a minimum acceptable offer before the time limit is reached, the computer will automatically reject the offer in that period. The following examples illustrate how earnings are calculated.

Example 1:
The sequence lasts for 14 periods. A participant's choice and the offers they draw are listed as follows:

Period	1	2	⋯	14
Minimum acceptable offer	243	-	⋯	-
Offer	675	-	⋯	-
Period Income	675	675	⋯	675

In Period 1, the offer of 675 is drawn and exceeds the chosen minimum acceptable offer, 243. Because this sequence lasts for 14 periods, the participant's earning from the sequence is 675*14 = 9,450 points.

Example 2:
The sequence lasts for 31 periods. A participant's choices and the offers they draw are listed as follows:

Period	1	2	3	4	5	6	7	⋯	31
Minimum acceptable offer	455	770	61	531	916	380	-	⋯	-
Offer	279	208	15	66	106	564	-	⋯	-
Period Income	0	0	0	0	0	564	564	⋯	564

In the first five periods none of the offers are accepted because they are less than the minimum acceptable offer. In Period 6, the offer of 564 is drawn and is accepted because it exceeds the minimum acceptable offer of 380. Because there are still 31-5=26 periods remaining in the sequence, the participant's earning from the sequence is 564*(31-5) = 14,664 points.

Example 3:
The sequence lasts for 8 periods. A participant's choices and the offers they draw are

listed as follows:

Period	1	2	3	4	5	6	7	8
Minimum acceptable offer	655	818	138	927	700	322	877	427
Offer	525	51	67	823	659	285	525	225
Period Income	0	0	0	0	0	0	0	0

None of the offers are accepted because they are all less than the minimum acceptable offer. Therefore the participant's earning from the sequence is 0 points.

The numbers in these examples have been chosen just for illustrative purposes. In the actual experiment, the computer will determine the number of periods in a sequence (using the rule that at the end of a period there is a 5% chance the sequence stops and a 95% chance the sequence continues). As long as you have not accepted an offer in the sequence, the computer will also draw the offer in each period (each offer from 1 to 1000 is equally likely).

[**Partner and Partner-Offer treatments only:** In each sequence, you will be randomly and anonymously paired with another participant. After the completion of each sequence, participants in the room will be randomly re-paired so the people you are paired with will change from sequence to sequence. In each sequence the person you are paired with is also choosing whether to accept offers. The offers you and that participant receive will be independent. This means you are equally likely to receive any offer from 1 to 1000, regardless of the other person's offer. When the computer determines whether to continue the sequence at the end of the period, it does this for both of you. That is, if the sequence continues, it continues for both of you and if it stops it stops for both of you.]

The minimum acceptable offer will be chosen by you in each period before an offer is accepted. [**Isolated treatment only:** At the end of each period, you will observe your offer and whether it is accepted (unless you already accepted an earlier offer). At the end of every period, you will receive feedback on period income. You will also observe your total income up to the period.] [**Partner treatment only:** At the end of every period, both you and the other participant will receive feedback on period income of both you and that participant. Both of you will also be informed of any offers made in that period and whether they were accepted. You will also observe your total income from the sequence up to the period.] [**Partner-Offer treatment only:** At the end of every period, both you and the other participant will receive feedback on period income

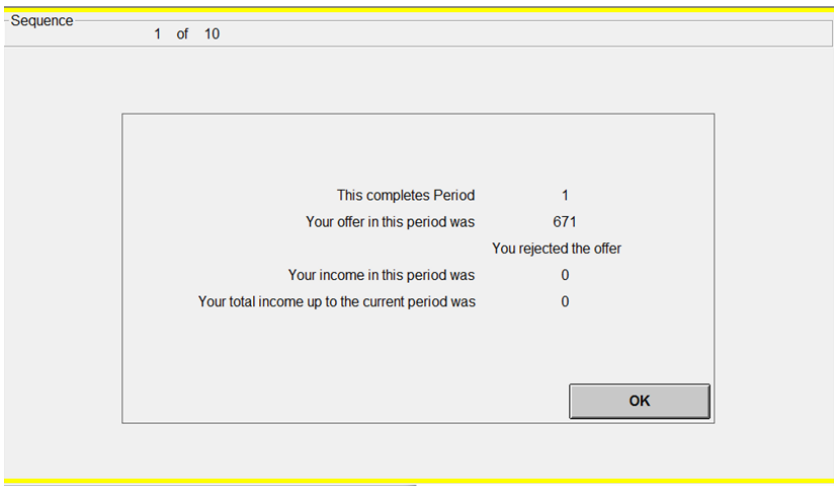
of both you and that participant. Both of you will also be informed of any offers made in that period and whether they were accepted. You will also observe your total income from the sequence up to the period.] The information will be displayed on the screen in the following layout:

[Screenshots: Figure 2.A1]

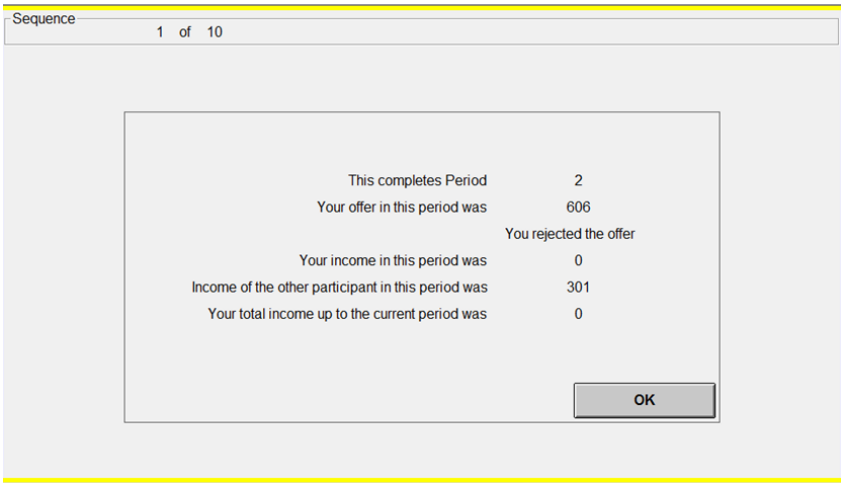
After the completion of each sequence, you will see the total number of periods, and your total income from all the periods in the sequence.

At the end of experiment, one of the 10 sequences will be drawn and the total points you earned in that sequence will be converted to cash as part of your payment.

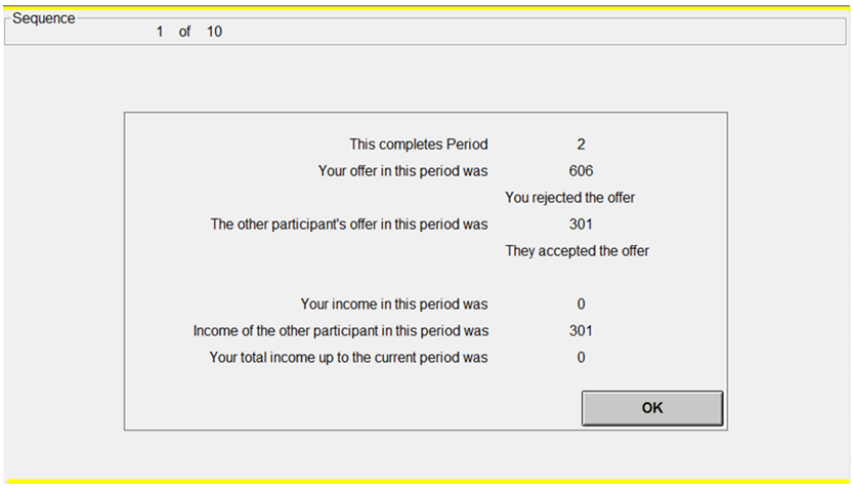
Before we begin the decision-making part of the experiment we will ask some questions to make sure you understand these instructions. Please look at you screen and follow the instructions. If at any time you have a question raise your hand and someone will come to your desk to answer it.



(a) Isolated Treatment



(b) Partner Treatment



(c) Partner-Offer Treatment

Fig. 2.A1 Screenshots for period feedback

Control questions

[*] *indicates differences across treatments.*

1. What is the probability that the computer will offer you more than 500 in Period 1 of a round?
- Lower than 50%
- Exactly 50%
- Higher than 50%

(Message if a wrong answer is selected: "The answer you picked is incorrect; the offer is drawn from 1, 2 to 1000 with equal probability, so in each period the probability of getting an offer larger than 500 is exactly 50%.")

2. Suppose in the Period 1 you rejected an offer of 233. What is the probability that the computer will offer you more than 500 in Period 2?
- Lower than 50%
- Exactly 50%
- Higher than 50%

(Message if a wrong answer is selected: "The answer you picked is incorrect; the offers you get are independently drawn, so the probability of getting an offer larger than 500 in Period 2 is exactly 50%.")

3. Suppose you accepted an offer of 879 in Period 1, what is the probability that the round will continue to Period 2?
- Lower than 95%
- Exactly 95%
- Higher than 95%

(Message if a wrong answer is selected: "The answer you picked is incorrect; the probability the round will continue to Period 2 is exactly 95%. At the end of every period there is a 5% chance the round stops and a 95% chance the round continues.")

4. Suppose you rejected an offer of 435 in Period 10, what is the probability that the round will continue to Period 11?
- Lower than 95%
- Exactly 95%
- Higher than 95%

(Message if a wrong answer is selected: "The answer you picked is incorrect; after Period 10, the probability that the round will continue to Period 11 is exactly 95%. At the end of every period there is a 5% chance the round stops and a 95% chance the round continues.")

5. In periods 1-3, the minimum acceptable offer and received offers of a participant are listed in the following table:

Period	1	2	3	···
Minimum acceptable offer	891	850	620	···
Offer	383	777	643	···

The participant has accepted the offer in which period?
- Period 1
- Period 2
- Period 3

(Message if a wrong answer is selected: "The answer you picked is incorrect; the participant has accepted the offer in Period 3, which is the first period that the offer received exceeds the minimum acceptable offer.")

[**Peer and Peer-Offer treatments only:**
6. Suppose in Period 1 your offer is 698, what is the probability that your paired participant's offer is larger than 500 in the same period?
- Lower than 50%
- Exactly 50%
- Higher than 50%

(Message if a wrong answer is selected: "The answer you picked is incorrect; the offers you and your paired participant get are independent, so for them the probability of getting a number larger than 500 is exactly 50%.")]

7. The experiment will consist of 10 sequences. After the completion of the experiment, how many sequence(s) will be drawn to determine the payoffs of the participants?
- Only one
- More than one
- The number of sequences is randomly determined by the computer

(Message if a wrong answer is selected: "The answer you picked is incorrect; only ONE out of 10 sequences will be drawn to determine the payoff of the participants. ")

Appendix 2.B Theoretical Analysis

2.B.1 Basic setup

We describe a simple job-search model based on McCall (1970). An infinitely-lived job-seeker makes decisions and receives income over a sequence of discrete time periods.

While searching the job-seeker draws one wage offer w per period from a bounded wage distribution $F(w)$. Offers are independently drawn across periods. If the job-seeker rejects the offer their income is zero in that period and they continue search. If the job-seeker accepts the offer, they exit search and receive w for the current and all future periods. In the baseline case we assume that the period utility is

$$u_t = y_t, \text{ where } y_t = \begin{cases} w \text{ if offer } w \text{ has been drawn and accepted,} \\ 0 \text{ otherwise.} \end{cases}$$

We assume the job-seeker maximises the expected discounted sum of period utilities:

$$\sum_{t=0}^{\infty} \delta^t E[u_t], \;\; 0 < \delta < 1$$

Since the environment is stationary, the optimal strategy can be expressed as a constant reservation wage r, i.e. in any period, the job-seeker accepts if offered at least the reservation wage, and continues to search otherwise. The job-seeker's objective function can be written as a function of the reservation wage, $V(r)$. With probability $F(r)$ the job-seeker rejects the offer, earns zero in the current period and faces an identical decision problem in the next period (and hence gets $\delta V(r)$). With probability $1-F(r)$ the job-seeker accepts the offer and earns the offered amount in all future periods. Thus,

$$V(r) = F(r)(0 + \delta V(r)) + (1 - F(r)) \frac{1}{1-\delta} E[w | w \geq r],$$

which can be re-arranged as

$$V(r) = \frac{1}{(1-\delta F(r))} \int_{w>r} \frac{w}{1-\delta}\, dF(w). \tag{2.B1}$$

For a given offer distribution this can be used to solve the optimization problem directly. For ease of exposition, in the remainder of the discussion we assume that the wage offer follows a uniform distribution over $[0,1]$, $F(w) = w$. Then (2.B1) is maximised with a unique r over $[0,1]$,

$$r^* = \frac{1-\sqrt{1-\delta^2}}{\delta}. \tag{2.B2}$$

For a discount factor $\delta = 0.95$, this implies an optimal reservation wage of $r^* = 0.724$. For our experiment with wage offers following a discrete uniform distribution on $\{1, \cdots, 1000\}$, the optimal reservation wage is 725.

2.B.2 Distributional preferences

Now suppose that there are two job-seekers, indexed by $i \in \{1,2\}$, who are "partners" in the sense that they perform the same sequential search task simultaneously, with wage offers being independent draws across job-seekers and periods. The job-seekers are informed of each other's income at the end of each period. We assume that the job-seekers make interpersonal comparisons in each period of search. Specifically, following Fehr and Schmidt (1999), we assume that i dislikes income differences, as represented by the period utility function (where y_{it} denotes i's period t income and y_{jt} denotes the period t income of i's partner)

$$u_{it}(y_{it}, y_{jt}) = y_{it} - \alpha\big(\max\{y_{it}, y_{jt}\} - y_{it}\big) - \beta\big(\max\{y_{it}, y_{jt}\} - y_{jt}\big) \tag{2.B3}$$

The assumptions made in Fehr and Schmidt (1999) are that $0 \leq \beta \leq 1$ and $\alpha \geq \beta$. We retain the assumption that utility is evaluated each period, including the distaste for inequality, and i seeks to maximise

$$\sum_{t=0}^{\infty} \delta^t E[u_{it}], \;\; 0 < \delta < 1 \tag{2.B4}$$

This extends the individual choice problem of Section 2.B.1 to a game.

Our solution concept is that of Markov Perfect Equilibrium. A Markov strategy for job-seeker i is $(r_{i0}, r_i(y))$, an initial reservation wage r_{i0} that governs i's acceptance decision when both partners are active job-seekers, and a reservation wage function $r_i(y)$ that governs i's acceptance decision when she is actively searching but the partner has already accepted an offer and obtains period income y. A Markov Perfect Equilibrium is a pair of Markov strategies $(r^*_{i0}, r^*_i(y))$, $i = 1, 2$, such that,

1. For each i, $r^*_i(y)$ maximises (2.B4) given partner j is inactive with period income y.

2. For each i, r^*_{i0} maximises (2.B4) given partner j is active using strategy $(r^*_{j0}, r^*_j(y))$, and i will revert to $r^*_i(y)$ if j exits with period income y.

We start by calculating $r^*_i(y)$ and then find r^*_{i0} using backward induction.

Optimal reservation wage for active job-seeker after partner exits

Consider the case that i is still searching and j has exited with a wage (and hence a per-period income) of $y_{jt} = y$. Again the dynamic problem facing i is stationary, and so the optimal reservation wage is constant.

Using a constant reservation wage r_i, i's expected discounted utility beginning from a period in which j is inactive with period income y is denoted by $V_A(r_i; y)$. With probability r_i the offer is rejected ($w < r_i$) , which gives the current period utility of

$-\alpha y$ and a discounted future value of $\delta V_A(r_i;y)$, and with probability $1-r_i$, the offer is accepted ($w \geq r_i$), which yields a constant period utility in all future periods. Thus,

$$V_A(r_i;y) = r_i(u_i(0,y) + \delta V_A(r_i;y)) + (1-r_i)\frac{1}{1-\delta}E[u_i(w,y)|w \geq r_i],$$

and rearranging gives

$$V_A(r_i;y) = \frac{1}{1-\delta r_i}\left(r_i u_i(0,y) + \frac{1}{1-\delta}\int_{r_i}^{1} u_i(w,y)\, dw\right).$$

With u_i given by (2.B3),

$$V_A(r_i;y) = \frac{1}{1-\delta r_i}\left(-r_i\alpha y + \frac{1}{1-\delta}\left(\int_{r_i}^{k}(w-\alpha(y-w))\, dw + \int_{k}^{1}(w-\beta(w-y))\, dw\right)\right),$$

where $k \equiv \max\{r_i, y\}$.

This is maximised by $r_i^*(y)$,

$$r_i^*(y) = \max\left\{\frac{1-\sqrt{1-\delta^2+2\delta(1-\delta)\,\frac{\alpha+\beta}{1-\beta}y}}{\delta}, \frac{1-\sqrt{1-\delta^2+\delta^2\,\frac{\alpha+\beta}{1+\alpha}(1-y)^2}}{\delta}\right\}. \tag{2.B5}$$

Equation (2.B5) describes the optimal reservation wage for a job-seeker (with given α, β and δ) after their partner has exited search with period income y. When $\alpha = \beta = 0$ this reduces to (2.B2) and the reservation wage is independent of the partner's income. For other values of α, β, $0 < r_i^*(y) \leq r^*$. For $y = 0$ or 1 the optimal reservation wage is r^*. For $y < \hat{y}$, $r_i^*(y)$ is decreasing in y. For $y > \hat{y}$, $r_i^*(y)$ is increasing in y.

We use $V_A(y)$ to denote the value of the game to i beginning from a period in which she is active but j has previously accepted an offer and receives period income y. That is,

$$V_A(y) = V_A(r_i^*(y);y). \tag{2.B6}$$

Optimal initial reservation wage when both job-seekers are active

Now we move on to find the equilibrium initial reservation wage. Job-seeker i's expected discounted utility beginning from a period in which both job-seekers are active is a function of their initial reservation wages $V_0(r_{i0}, r_{j0})$. The period has four possible outcomes: neither of the job-seekers accept their offers, i accepts and j rejects, j accepts and i rejects, or both of the job-seekers accept. Thus $V_0(r_{i0}, r_{j0})$ can be written as the sum of four components, each regarding the payoff from one of the outcomes:

$$\begin{aligned} V_0(r_{i0}, r_{j0}) = &\int_0^{r_{j0}} \int_0^{r_{i0}} \left(0 + \delta V_0(r_{i0}, r_{j0})\right)\, dy_j dy_i \\ &+ \int_0^{r_{j0}} \int_{r_{i0}}^{1} \left(u_i(y_i, 0) + \delta V_I(y_i)\right)\, dy_i dy_j \\ &+ \int_0^{r_{i0}} \int_{r_{j0}}^{1} \left(u_i(0, y_j) + \delta V_A(y_j)\right)\, dy_j dy_i \\ &+ \frac{1}{1-\delta} \int_{r_{i0}}^{1} \int_{r_{j0}}^{1} u_i(y_i, y_j)\, dy_j dy_i \end{aligned}$$

When neither accepts, i receives current period utility of zero and in the next period begins a subgame which is identical to the original game. When i accepts and j rejects, i's current period utility is $u_i(y_i, 0)$, and we denote by $V_I(y_i)$ (subscript I for "inactive") the value of the game to i beginning from a period in which she is inactive with period income y_i and the partner is active. When i rejects and j accepts and earns period income y_j, i's current period utility is $u_i(0, y_j)$, and the value of the game to i from the beginning of the next period is $V_A(y_j)$. Finally, when both i and j accept their offers in the current period i gets utility $u_i(y_i, y_j)$ in this and every future period.

Rearranging and using (2.B3) gives

$$V_0(r_{i0};r_{j0}) = \frac{1}{1-\delta r_{i0}r_{j0}}\Bigg(r_{j0}\int_{r_{i0}}^{1}\big((1-\beta)y_i+\delta V_I(y_i)\big)\,dy_i + r_{i0}\int_{r_{j0}}^{1}\big(-\alpha y_j+\delta V_A(y_j)\big)dy_j \\ +\frac{1}{1-\delta}\int_{r_{i0}}^{1}\int_{r_{j0}}^{1}\big(y_i-\alpha\big(\max\{y_i,y_j\}-y_i\big)-\beta\big(\max\{y_i,y_j\}-y_j\big)\big)\,dy_j dy_i\Bigg), \tag{2.B7}$$

Note that $V_I(y_i)$ depends on both job-seeker's preferences (and i's belief about j's preferences), so identifying the equilibrium initial reservation wage requires assumptions about both job-seekers' preferences and what job-seekers know about each other's preferences. For tractability, we assume job-seekers have identical preferences and this is common knowledge.

If i is inactive with period income y_i we assume j uses the optimal reservation wage $r_j^*(y_i)$. Then,

$$V_I(y_i) = r_j^*(y_i)\big(u_i(y_i,0)+\delta V_I(y_i)\big)+(1-r_j^*(y_i))\frac{1}{1-\delta}E[u_i(y_i,w)|w\geq r_j^*(y_i)],$$

which is rearranged as

$$V_I(y_i) = \frac{1}{1-\delta r_j^*(y_i)}\Bigg(r_j^*(y_i)u_i(y_i,0)+\frac{1}{1-\delta}\int_{r_j^*(y_i)}u_i(y_i,w)\,dw\Bigg).$$

Given (2.B3), it is convenient to denote $l\equiv\max\{r_i^*(y_i),y_i\}$, so we have

$$V_I(y_i) = \frac{1}{1-\delta r_i^*(y_i)}\Bigg(r_i^*(y_i)(1-\beta)y_i \\ +\frac{1}{1-\delta}\big((1-r_i^*(y_i))y_i-\alpha\int_l^1(w-y_i)dw-\beta\int_{r_i^*(y_i)}^{l}(y_i-w)dw\big)\Bigg). \tag{2.B8}$$

Combining (2.B5), (2.B6), (2.B7) and (2.B8), we complete the initial objective function V_0 for i. We conduct a calibration to find the optimal response function for

i given j's choice, $r^*_{i0}(r_{j0})$. We calibrate r^*_{i0} for all (α,β), $\alpha \in [0,2)$, $\beta \in [0,1)$, and $\alpha > \beta$ with grid equal to 0.01. The results are accurate to the third decimal place.

The simulation shows that the equilibrium r^*_{i0} is close to the expected own-earning-maximizing r^*. For the majority of (α,β), r^*_0 is slightly below r^*, although for small α and small β it is slightly above r^* (Figure 2.B1a).

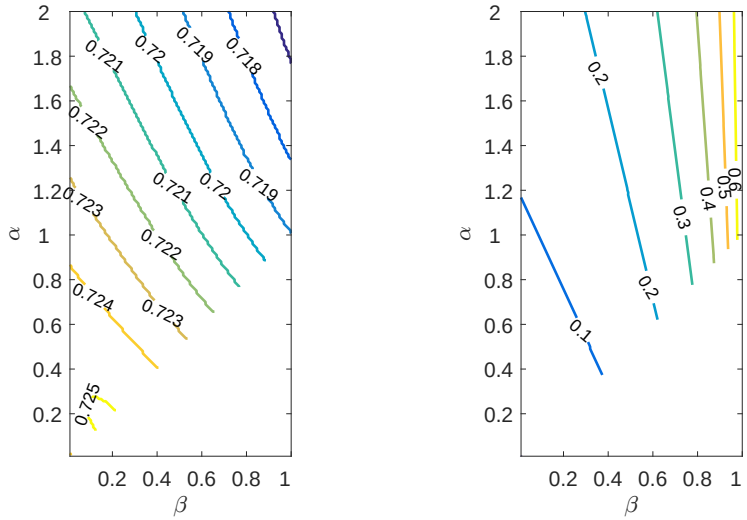
The initial equilibrium reservation wage r^*_{i0} lies above the kink point of the $r^*_i(y)$ function, $\hat{y}$. This implies that a job-seeker never accepts a wage below $\hat{y}$ (Figure 2.B1b) and so in equilibrium only the positively-sloped part of $r^*_i(\cdot)$ is relevant in Figure 2.1.

Summary and implications

After a partner exits, the model predicts that the remaining job-seeker selects a constant reservation wage that depends on the partner's income (see Figure 2.1). Since the initial reservation wage lies above the kink point, only the upward sloping portion is relevant in equilibrium. Therefore the model predicts a positive relationship between the wage accepted by the partner and the reservation wage of the remaining job-seeker. In most cases, the change in reservation wage in response to partner's exit is negative, but the size of the change is moderated by the exit offer of the partner. Note that although we only report the calibration with standard Fehr-Schmidt assumptions, this qualitative prediction can be extended to negative β's, as long as $-\alpha < \beta$.

These results have three practical implications for job search outcomes. First, the probability that a job-seeker exits in period t depends positively on whether their partner exited in period $t-1$. In our setting, the probability of exiting is $1-F(r) = 1-r$, which changes from $(1-r^*_0)$ before the partner exits to $(1-r^*(y))$ after the partner exits with y. The expected change is $r^*_0 - E[r^*(y)|y \geq r^*_0]$. Before the partner exits, the probability of exiting is about 28%, and the expected increase in the probability of exiting after

Fig. 2.B1 Equilibrium initial reservation wage for different Fehr-Schmidt parameters



(a) Initial equilibrium reservation wage r_0^*

(b) Comparison of r_0^* with kink-point: $r_0^* - \hat{y}$

Notes: $\beta \in [0,1)$, $\alpha \in [0,2)$, $\alpha \geq \beta$; $r^* = 0.724$. The equilibrium r_{i0}^* lies between 0.717 and 0.726 for all (α,β). In the calibrated range of (α,β), r_{i0}^* takes its minimum at $(1.99,0.91)$ and maximum at $(0.24,0.06)$. $(r_{i0}^* - \hat{y})$ is always positive, and takes its minimum of 0.003 at $(0.01,0.01)$ and maximum of 0.687 at $(1.99,0.99)$.

partner exits varies from about 3% for high values of α and β to 1% for low values. The expected changes for five sets of (α,β) are reported in column (1) of Table 2.B1.

Second, there is a "spillover" from the partner's exit wage to the remaining job-seeker's reservation wage. A job-seeker whose partner has exited with a high wage has a higher reservation wage than a job-seeker whose partner has exited with a low wage. Given an initial reservation wage of r_0^*, the exit wage varies uniformly between r_0^* and 1. The reservation wage of the active job-seeker then varies from $r^*(r_0^*)$ to $r^*(1)$. The difference between $r^*(1)$ and $r^*(r_0^*)$ varies from 0.10 for high α and β to 0.04 for low α and β. Column (2) of Table 2.B1 reports the slope of the linearised relationship between the partner's exit wage and the remaining job-seeker's reservation wage, which shows

Table 2.B1 Quantifying the implications for search outcomes

	(1)	(2)	(3)
	Expected change in prob. of exiting when j exits:	Linearised relationship between. $r^*(y)$ and y, $y \geq r_0^*$:	Correlation between exit wages:
$(\alpha,\ \beta)$	$r_0^* - E[r(y)\|y \geq r_0^*]$	$\frac{r^*(1)-r^*(r_0^*)}{1-r_0^*}$	$\rho(y_1, y_2)$[a]
(0, 0)	0.000	0.000	0.000
(0.5, 0.0)	0.013	0.132	0.046
(0.5, 0.5)	0.024	0.253	0.100
(0.9, 0.6)	0.027	0.297	0.110
(1.0, 0.8)	0.029	0.336	0.129
(2.0, 0.8)	0.028	0.349	0.130

[a] Based on calibration of 100,000 pairs.

that a unit increase in y leads to a linearised increase in r^* of 0.35 for high and 0.13 for low α and β.

Third, this implies a correlation of wages across partners. By definition, independent job-seekers' wages are uncorrelated. In contrast, inequity averse job-seekers who are partners will have positively correlated wages, if they exit in sequence. The correlation is 0.13 for high α and β and 0.05 for low α and β (column (3) of Table 2.B1).

Appendix 2.C Additional results

Table 2.C1 Alternative specifications for Table 2.4 on between-period change in reservation wage

	(1) Isolated			(2) Partner			(3) Partner-Offer		
β_1 j accepted offer (a_j)	−11.443	−8.923	−16.231	−71.029**	−69.019*	−71.662**	−58.712***	−61.161***	−59.495***
	(11.988)	(13.411)	(12.675)	(33.712)	(38.390)	(35.521)	(17.223)	(16.297)	(19.022)
β_2 j's income ($a_j o_j$)	0.021	0.018	0.028	0.094**	0.088*	0.092*	0.083***	0.082***	0.083***
	(0.027)	(0.032)	(0.027)	(0.047)	(0.052)	(0.049)	(0.025)	(0.022)	(0.029)
β_3 value of j's offer (o_j)	0.004	0.003	0.003	−0.008	−0.001	−0.006	−0.002	−0.001	−0.002
	(0.016)	(0.019)	(0.016)	(0.022)	(0.022)	(0.023)	(0.012)	(0.011)	(0.012)
Subject random effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Period fixed effects	No	No	Yes	No	No	Yes	No	No	Yes
Sequence fixed effects	Yes	No	Yes	Yes	No	Yes	Yes	No	Yes
Subject characteristics	No	No	Yes	No	No	Yes	No	No	Yes
Number of obs.	502	502	502	578	578	578	591	591	591
Number of sessions	6	6	6	6	6	6	6	6	6

Table 2.C2 Estimates of sequence trend

	Period 1 (r_{is1}) (1)	(2)	Exit period (r_{isL}) (3)	(4)
Isolated	3.733	3.732	4.078***	4.083***
	(2.555)	(2.559)	(1.293)	(1.294)
Partner	9.538***	9.540***	12.926***	12.904***
	(1.883)	(1.886)	(2.235)	(2.239)
Partner-Offer	6.333**	6.328**	7.663***	7.659***
	(2.875)	(2.877)	(2.463)	(2.466)
p-values for				
Partner vs. Isolated	0.067	0.068	0.001	0.001
Partner-Offer vs. Isolated	0.499	0.500	0.197	0.199
Partner vs. Partner-Offer	0.351	0.350	0.113	0.115
Subject random effects	Yes	Yes	Yes	Yes
Subject characteristics	No	Yes	No	Yes
Number of obs.	2,155	2,155	2,035	2,035

Notes: The sequence trends for each treatment are estimated using the following model:
$r_{ist} = \beta_0 + \beta_1 \text{Partner}_i + \beta_2 \text{Partner-Offer}_i + \beta_3 s + \beta_4 s \times \text{Partner}_i + \beta_5 s \times \text{Partner-Offer}_i + \eta_i + \varepsilon_{ist}$,
where t is period 1 or the last period. The sample for column (1) and (2) excludes 5 sequences in which the subject timed out for the initial period. The sample for column (3) and (4) excludes 135 censored sequences.

Chapter 3

Meeting and matching:

Decomposing the matching function

[*] This chapter is based on a joint work with Martyn Andrews and Richard Upward.

3.1 Introduction

There is a large empirical literature which has estimated and tested various components of the search and matching framework. In this paper we directly estimate two components of the matching function: uniquely, we use micro-level data to decompose both vacancy and job-seeker matching rates into the matching probability and the arrival rate of interview meetings from both sides of the same well-defined labour market.

Decomposing the matching rate into its constituent parts allows us to answer three questions which are central to the search and matching literature.

First, there is a large empirical literature which assumes and estimates a Cobb-Douglas matching function which relates the number of matches to stocks of job-seekers (U) and vacancies (V). Typically, these estimates suggest that there are constant returns to scale in U and V. But, as noted by Eckstein and van den Berg (2007), parameter estimates from a reduced-form hazard model cannot separate between the relative magnitudes of the job offer arrival rate and the acceptance probability, or estimate the relative magnitudes of the effects of the variables on them. Our estimates allow us to calculate precisely these magnitudes. For example, does the effect of labour market tightness operate through the meeting rate or the matching probability? In two-sided search models one expects that the effect of U on the job-seeker's reservation wage is canceled out by the effect of U on the employer's reservation profit for the vacancy, in which case the combined matching probability will not be a function of U or V (Burdett and Coles, 1999).

Second, we can examine whether changes in the per-period matching probability (or hazard rate) over the search duration is determined by changes in the arrival rate of suitable partners, or changes in the matching probability conditional on arrival. In the raw data, both job-seekers and vacancies tend to have declining hazard rates as search duration increases. Our decomposition reveals whether this is because suitable

partners become harder to contact as time passes, or because the matching probability declines. There is also a statistical benefit from decomposing the hazard rate into a meeting hazard and a matching hazard. Estimates of the meeting hazard can be identified using repeated observations on the same agents, in contrast to estimates of the matching hazard. This allows us to address directly the issue that duration dependence may be caused by unobserved heterogeneity.

Third, our agent-level data contains various characteristics of job-seekers (including educational qualifications, age and sex) and vacancies (including occupation, skill-level, firm size). These characteristics may be correlated with labour market tightness, and therefore we can test whether their inclusion changes our estimates of the parameters of interest. Furthermore, the decomposition allows us to assess whether characteristics which affect the matching rate do so because of their effect on the meeting rate or the matching probability.

Our central aim in this paper is to explore the decomposition of the matching rate into its constituent parts. In addition, our data also shed light on recurring issues in the empirical matching literature. It is well-known that data recorded at discrete frequency (e.g. monthly) may lead to biased estimates of the parameters of the matching process which is instantaneous (Berman, 1997; Burdett et al., 1994; Coles and Petrongolo, 2008; Gregg and Petrongolo, 2005). It has also been discussed in the literature that spatial aggregation can lead to biases: Coles and Smith (1996) discuss the problem with aggregating separate labour markets as a single market, and other works (Burda and Profit, 1996; Burgess and Profit, 2001; Patacchini and Zenou, 2007; Petrongolo and Manning, 2017) demonstrate the importance of accounting for interdependence between neighbouring labour markets.

Our use of micro data from both sides of the same market shows that misspecification is likely to be a problem in many estimates of the matching function. In particular,

temporal aggregation bias and cross-district search mean that a naive estimate of the matching function produces biased estimates, which are divergent from the two sides of the market. We show that, using a specification which incorporates suggestions of both strands of the literature (i.e. accounting for the spatial spill-overs, and correcting for the temporal aggregation bias), the estimates of the parameters of the matching function move closer together.

Our findings are as follows. Constant returns to scale are not rejected for the estimates of the matching function from both sides of the market. There is a mild increasing return to scale for the meeting function, due to a smaller congestion effect among job-seekers in terms of getting a meeting compared to getting a match. The labour market tightness mainly takes effect through the meeting decision, not through the matching probability. This finding is qualitatively consistent with the findings for Berman (1997) using Israeli data.

We find that the meeting rate exhibits a strong negative duration dependence for both job-seeker and vacancy after controlling for observed characteristics and unobserved heterogeneity, and no strong duration effect is found in the matching probability. Discrimination against long-term unemployed could explain the findings on the job-seeker side, but not on the vacancy side. Whereas other explanations of duration dependence, such as stock-flow matching, could consistently explain a (symmetric) duration dependence on both sides of the market.

In terms of characteristics, the effects on contact rates and matching probabilities are often but not always in the opposite direction. For example, job-seekers with a high education qualification have higher matching hazards, which arises from much higher contact hazards and a slightly lower matching probability. High-quality vacancies, however, have lower contact hazards and lower matching hazards.

The rest of the paper is structured as follows. In Section 3.2 we discuss the related empirical literature. In Section 3.3 we present a simple statistical model for the decomposition, and explain how we deal with the temporal aggregation bias and geographic spill-overs. We describe the data in Section 3.4. We present the results in Section 3.5, and Section 3.6 offers a discussion and some conclusions.

3.2 Literature

The idea that the matching rate can be decomposed into a meeting rate and a matching probability originates from Pissarides (1990). If job-seekers and vacancies are homogenous then every meeting will result in a match. But, if job-seekers and vacancies are heterogenous, and the productivity of a pair can only be revealed after meeting, then not every meeting will result in a match. When a pair meets, if the match productivity is high enough in relation to alternatives for both the job-seeker and the employer the pair will match and otherwise they will go back to search. Therefore the matching rate is given by the product of the meeting rate and the probability that both partners are acceptable to the other. The effect of labour market tightness ($\theta = V/U$) on the matching rate can be decomposed into a "meeting" effect (e.g. an increase in θ increases job-seekers' contact rate) and a "reservation wage" effect (e.g. an increase in θ makes job-seekers more selective). In this section we discuss studies which disentangle the two effects. As noted in the introduction, the decomposition of matching rates into their constituent parts is related to three different areas of the search and matching literature.

3.2.1 Estimates of the matching function and its components

In the current study, we are able to decompose the matching rate without relying on any structural assumptions because we directly observe meetings and matches. Berman

(1997) is the only previous study which is able to do this decomposition directly. Berman (1997) uses the data from the Israeli Employment Service which documents the aggregate time series of referral meetings, matches, and stocks of job-seekers and vacancies. This allows Berman to estimate aggregate meetings using a Cobb-Douglas specification after addressing data issues which we discuss in Section 3.3. Berman finds that market tightness affects the hiring rate through affecting the meeting rate, for which constant returns to scale are not rejected, and that there is only small and insignificant effects of market conditions on the probability of a referral becoming a hire. In addition, he reports that the inflow stocks of job-seekers and vacancies generate more meetings than the incumbent stocks, which is interpreted as evidence of a negative duration dependence for both job-seekers and vacancies.

In other studies, since meetings are not separately observed, partial equilibrium models are applied to facilitate the decomposition. Given the assumptions about how a specific agent or labour-market characteristic influences the meeting technology and/or the productivity distribution, the model for optimal search decision generates predictions about how this characteristic influences the hazard rate. The structural estimation of the model only requires data of unemployment durations, reservation wages or information of rejected job offers, and income variables (e.g. Blau and Robins, 1986; Bloemen and Stancanelli, 2001; Cobb-Clark et al., 2004; Frijters and van der Klaauw, 2006; García-Pérez, 2006; Jensen and Westergård-Nielsen, 1987; Lollivier and Rioux, 2010; Petrongolo and Pissarides, 2006; Stern, 1989; van den Berg, 1990).

3.2.2 Duration dependence of hazard rates

It is a well-documented fact that exit rates from unemployment decreases for those who are unemployed for longer. However, there is no universal answer to whether the observed duration dependence is caused by sample selection (Machin and Manning,

1999), or is a genuine effect, i.e. the exit rate decreases over time for anyone who is unlucky. The true duration dependence can be caused by the matching mechanism (e.g. the stock-flow matching model of Coles and Smith, 1998) or the deterioration of search perspective (combined with the assumption that the long-term unemployed do not sufficiently lower their reservation wage), changes in characteristics such as human capital (e.g. Arulampalam, 2001) and networks (e.g. Bramoullé and Saint-Paul, 2010), or discrimination (e.g. Jarosch and Pilossoph, 2015; Kroft et al., 2013).

Using micro data, studies can control for unobserved heterogeneity in analyzing duration dependence. Using UK data, Arulampalam (2001) find a strong state dependence, i.e. those who have been unemployed in the previous year are more likely to be unemployed, which could be a combination of incidence or duration dependence. Using French data, Abbring et al. (2002) find that the negative duration dependence is negligible at shorter durations, whereas at longer durations there is negative duration dependence. Using Danish data with repeated unemployment spells for some individuals, Rosholm (2000) also reports a negative duration dependence, although the finding only applies to some of the age and experience groups.

Several studies shed light on the duration dependence of offer arrivals (which result from meetings) and the offer distribution (which, along with the reservation wage, determines the matching probability conditional on the meeting). In Spanish data, García-Pérez (2006) finds that the unemployment hazard is approximately constant for the first six months because, although the job offer arrival rate is falling, the acceptance rate rises as the reservation wage falls. However, after six months the acceptance rate is close to one, and so the fall in the unemployment hazard is driven entirely by the fall in the offer arrival rate. However, using French data, Lollivier and Rioux (2010) find that negative duration dependence in the unemployment hazard rate arises from a deterioration in the value of wage offers, causing a decline in the acceptance rate.

Similarly, using German data, Frijters and van der Klaauw (2006) also find that the value of wage offers declines with search duration, while the offer arrival rate is constant.

There are relatively little evidence on the duration dependence of vacancies and the findings are again mixed. Using micro data from Netherlands, van Ours and Ridder (1992) report a positive duration dependence for vacancies. Using US data, Barron et al. (1999) find that vacancies tend to have very short durations - in one market most filled vacancies are filled immediately, leading to a negative duration dependence. Using the Lancashire Matching Service data, Andrews et al. (2008) find that hazards to filling exhibit negative duration dependence, whereas hazards to withdrawal exhibit positive duration dependence.

3.2.3 Effects of agent characteristics on hazard rates

The impact of job-seeker characteristics, decomposed into the impact on offer arrival rate and on the wage offer distribution has been investigated in Petrongolo and Pissarides (2006). Using micro survey data, they estimate the job-offer arrival rate and the mean wage offer as a function of individual and market-level characteristics. They find that labour market size affects the hazard rate through the offer distribution, not the offer arrival rate. On the job-seeker characteristics they report that male, skilled and older workers sample wage offers from a distribution with higher mean than female, unskilled and younger workers, and the job-offer arrival rates are higher for female, skilled, younger workers than for male, unskilled and older workers.

The impact of job-seeker characteristics on the probability of match after an application or interview has been investigated in several previous studies. Using French local employment office data, Teyssière (1995) finds that employers discriminate against the youngest applicants, the long-term unemployed and ethnic minorities. In a database of filled vacancies and applications to those vacancies, Russo and van Ommeren (1998)

find that the probability of a vacancy being filled by a female is proportionate to the share of female applicants. This suggests that female applicants not discriminated against, although this could result from self-selection in the applications. Using the Lancashire Matching Service data, Andrews et al. (2001) also find no discrimination against females, but there is evidence of discrimination against ethnic minorities.[1]

Cleaner evidence of discrimination is provided by field experiments using audit and correspondence methods, i.e. investigating the call-back and/or offer rates for pairs of fictitious resumes and/or candidates differ in only one characteristic (reviewed in Bertrand and Duflo, 2017). Kroft et al. (2013) find that the call-back rate is significantly lower for the long-term unemployed, and such discrimination is weaker in the slack labour markets, which suggests that the employers recognize that the duration signal is less informative in recessions. This literature tends to confirm that there is a strong discrimination based on ethnic background, but discrimination against women is less apparent (see Riach and Rich, 2006, for UK evidence).

On the employer characteristics, van Ours and Ridder (1991) find that job-filling rate is lower for vacancies which require higher educational level. Davis et al. (2013) find that job-filling rates decline with employer size, falling by more than half in moving from small to large establishments.

There three distinctions between our study and the existing literature. First, since we observe both sides of the market, we can test our identification of the meeting and matching functions to achieve an unbiased estimate. Second, since we observe agent-level data, we can address the source of duration dependence more directly and accurately. And third, we estimate the effects of both job-seeker and vacancy characteristics in a unified framework.

[1]The contacts examined in this paper is the subset of the those examined in Andrews et al. (2001), as we only focus on the unemployed job-seekers and job vacancies, and exclude contacts which involve job-seekers in education or employment, and vacancies in youth training programs.

3.3 Econometric Methods

3.3.1 Decomposing job-seeker and vacancy hazards

In this subsection, we set out a generic two-sided model of search in a stochastic environment. We do this is to show how the matching hazards can be decomposed, from a statistical perspective, into a contact hazard and a matching probability. This is done for both sides of the market and does not depend on the economic model by which search and matching takes place.

In any search and matching model, there are stocks of unmatched vacancies, denoted V, and unmatched job-seekers, denoted U, attempting to contact each other on a particular day in a particular market. This is modelled by the so-called contact function (sometimes known as the encounter function). Seen from the point of view of a single job-seeker w, the number of vacancies she encounters every day, denoted C^w, is Poisson distributed:

$$C^w \sim \text{Poisson}(\lambda^w). \tag{3.1}$$

λ^w is the Poisson parameter that denotes the daily average number of contacts per job-seeker. We denote the observed number of contacts for i-th job-seeker as C_i, a non-negative integer.

Every time a vacancy is encountered, the pair either consummates the contact by matching, or they do not. Agents from unsuccessful contacts return to the stocks of U and V. The probability that a given contact results in a match is denoted by μ,

$$M|C \sim \text{Bernoulli}(\mu), \tag{3.2}$$

where M denotes whether or not a particular contact results in a match. The observed outcome is either $M_i = 1$ if the job-seeker matches or $M_i = 0$ if she does not. Once

a job-seeker has matched successfully, there are no more matches. It follows that the marginal distribution for the number of job-seeker matches per period also has a Poisson distribution:

$$M^w \sim \text{Poisson}(\lambda^w \mu). \tag{3.3}$$

In words, if λ^w vacancies are encountered on average each day and a proportion μ of them match on average, then the average daily matching rate of job-seekers is $\lambda^w \mu \equiv h^w$. h^w is the job-seeker hazard. The marginal distribution for M^w shows that, if one had data on matches only, one cannot separately identify μ and λ^w.

All of the above is repeated for the other side of the market. In particular:

$$C^e \sim \text{Poisson}(\lambda^e), \qquad M^e \sim \text{Poisson}(\lambda^e \mu),$$

where λ^e is a Poisson parameter that denotes the average arrival rate of job-seekers to a given vacancy (labelled e). This means that the average daily matching rate of vacancies, or the vacancy hazard, is $\lambda^e \mu \equiv h^e$.

It is assumed that pairs of agents are drawn *randomly* from U and V and contact each other according to the Poisson process just described. On average, the total number of matches in the market multiplies h^w by the total number of job-seekers U, and this aggregate matching function is specified as a Cobb-Douglas function:

$$h^w U = e^{\kappa} U^{\alpha} V^{\beta}.$$

Hence the hazard for job-seekers is

$$h^w \equiv \lambda^w \mu = e^{\kappa} U^{\alpha-1} V^{\beta}, \tag{3.4}$$

where we can decompose h^w as follows:

$$\lambda^w = e^{\kappa_1} U^{\alpha_1 - 1} V^{\beta_1}, \tag{3.5}$$

$$\mu = e^{\kappa_2} U^{\alpha_2} V^{\beta_2}. \tag{3.6}$$

Note that

$$\alpha = \alpha_1 + \alpha_2, \qquad \beta = \beta_1 + \beta_2, \qquad \kappa = \kappa_1 + \kappa_2. \tag{3.7}$$

Typically one expects that the matching function exhibits constant returns to scale, namely $\alpha + \beta = 1$. If the contact function does the same, $\alpha_1 + \beta_1 = 1$, then it follows that the matching probability should be homogenous of degree zero, namely $\alpha_2 + \beta_2 = 0$.

Seen from the employer side of the market, the aggregate number of matches is the same:

$$h^e V = e^{\kappa} U^{\alpha} V^{\beta}.$$

Hence, the hazard for vacancies is written

$$h^e \equiv \lambda^e \mu = e^{\kappa} U^{\alpha} V^{\beta - 1}, \tag{3.8}$$

where again we can decompose h^e into

$$\lambda^e = e^{\kappa_1} U^{\alpha_1} V^{\beta_1 - 1}, \tag{3.9}$$

and Equation (3.6). Note that it is the *same* set of Cobb-Douglas parameters $\{\alpha,\ \beta,\ \alpha_1,\ \beta_1\}$ that appear in the two vacancy equations (3.8) and (3.9) as in the two job-seeker models (3.4) and (3.5) above. This means that if the job-seeker data and vacancy data come from the same single labour market, the estimates of these parameters should be

the same from both sides of the market. Our data *are* from the same market, a feature that makes this paper unique.

3.3.2 Count data and binary choice estimation

A limitation of our data (and the data used in most previous studies) is that the events of interest are not observed daily. This poses a challenge for us (and others investigators) to estimate the rates of the events, which are determined by the daily U and V. More specifically, for contacts we do not observe the dates, only the months that they occur. Distinct from most of the literature, however, we do observe the dates when a job-seeker becomes unemployed/a vacancy is posted, when the job-seeker starts working if matched and when agents become inactive. This means we can approximate the date of match and measure the search duration in days instead of months. In order to model contacts, we organise the data into an unbalanced count-data panel of contacts and matches, where the unit of observation is a month of search (i.e. elapsed month) of a job-seeker/vacancy.

Our job-seeker panel data for job-seeker i in the n'th elapsed month is given by $\{C_{in}, M_{in}, \tau_{in}, x'_{in}\}$, where τ_{in} records the job-seeker's duration in the elapsed month and x'_{in} represents all other job-seeker covariates. Similarly, $\{C_{jn}, M_{jn}, \tau_{jn}, x'_{jn}\}$ describes the vacancy panel, including vacancy's duration τ_{jn} and covariates x'_{jn}. There is a third dataset, whose unit of observation is every contact $q_{(i,j)}$, and records whether or not the given contact results in a match $m_{q(i,j)} \in \{0,1\}$, shortened to m_{ij}. The contact data also includes covariates for both agents, and is described as $\{m_{ij}, x'_{ij}\}$.

Our econometric methods are as follows. From Equation (3.1), λ^w can be estimated using count-data methods on the job-seeker data just described. Similarly, From Equation (3.3), $h^w(\equiv \lambda^w \mu)$ can be estimated using the same data. From Equation (3.2), it is clear that μ can be estimated using any binary outcome methods with the contact-level

data. Exactly the same considerations apply to the vacancy data when estimating λ^e and h^e using vacancy data.

We now write down exactly what each regression comprises. To do this, note that the canonical Poisson link function is Exponential: $\lambda^w = \exp(x'\delta_1)$ and $h^w \equiv \lambda^w \mu = \exp(x'\delta)$. We start with Equations (3.4) and (3.5), and take logs:

$$\log h_{in}^w = \kappa + (\alpha - 1)\log U_{kt} + \beta \log V_{kt} + \log \tau_{in} + \theta_n^w + \text{other covs} + \varepsilon_i^w \tag{3.10}$$

$$\log \lambda_{in}^w = \kappa_1 + (\alpha_1 - 1)\log U_{kt} + \beta_1 \log V_{kt} + \log \tau_{in} + \theta_{1n}^w + \text{other covs} + \xi_i^w. \tag{3.11}$$

The stocks of job-seekers U_{kt} and vacancies V_{kt} are defined at the level of the local labour market k and vary by calendar month t. They are measured at the beginning of each month. The logged effect of job-seeker's duration in the elapsed month τ_{in} is constrained to unity. The θ_n^w terms represent a set of elapsed month dummies. These estimate a non-parametric baseline matching hazard. Similarly, the θ_{1n}^w estimate a non-parametric baseline contact rate.

The term "other covs" refers to all other relevant observable characteristics of job-seekers or vacancies. Notice that it is the same set of covariates for both (3.10) and (3.11). Moreover both sets of agent-level characteristics are averaged to a market–month kt, and are included in *both* job-seeker and vacancy regressions, to control for the differences in the quality of stocks for different local labour markets. All regressions include dummies for local labour market k, four quarter dummies, five year dummies, and five year dummies interacted with local labour market k. It is important to include these dummies to account for differences in matching technology over time and for different local labour markets, otherwise these aggregated shocks will lead to serial correlation in the residuals.

Finally, ξ_i^w and ε_i^w are job-seeker random effects. They are assumed to be uncorrelated with all covariates, but especially the dummies that model elapsed duration. This means we estimate the parameters of Equations (3.10) and (3.11) by a Poisson Random Effects estimator. These methods are routine and standard (see Wooldridge, 2010, Section 18.7.3), for the key assumptions and the resulting properties. We can also estimate Equation (3.11) by a Poisson Fixed Effects Estimator, and hence can test whether the ξ_i^w are uncorrelated with elapsed duration. This is because some job-seekers experience multiple months of contacts. As always, these estimates come from an unrepresentative sample because those job-seekers who match in their first month, or fail to have any contacts, are dropped from the estimation. This is why Poisson Fixed Effects estimates for Equation (3.10) are not reported because they are identified from only those job-seekers who experience multiple spells, which is a very small fraction of the sample.

Subtracting (3.11) from (3.10) generates estimates of α_2 and β_2, given Equation (3.7) above. This is also a model in its own right:

$$\log \mu_{ij} = \kappa_2 + \alpha_2 \log U_{kt} + \beta_2 \log V_{kt} + \theta_{2n}^w + \text{other covs} + (\varepsilon_i^w - \xi_i^w), \tag{3.12}$$

and can be estimated by the contact-level data. Whilst we could use any Random Effects Binary Choice estimator, again we use the Poisson because its canonical link is exponential. The θ_{2n}^w estimate a non-parametric baseline hazard for the matching probability. However, note that the exposure term is absent because the concept is meaningless in a dataset of contacts.

For completeness, the analogous vacancy regressions are written as:

$$\log h_{jn} = \kappa + \alpha \log U_{kt} + (\beta - 1) \log V_{kt} + \log \tau_{jn} + \theta^e_n + \text{other covs} + \varepsilon^e_j \tag{3.13}$$

$$\log \lambda^e_{jn} = \kappa_1 + \alpha_1 \log U_{kt} + (\beta_1 - 1) \log V_{kt} + \log \tau_{jn} + \theta^e_{1n} + \text{other covs} + \xi^e_j \tag{3.14}$$

$$\log \mu_{ij} = \kappa_2 + \alpha_2 \log U_{kt} + \beta_2 \log V_{kt} + \theta^e_{2n} + \text{other covs} + (\varepsilon^e_j - \xi^e_j). \tag{3.15}$$

Notice that we estimate a second model for μ, this time with vacancy covariates. Equation (3.10)–(3.15) define our base model. In the next two subsections we will introduce corrections to this base specification. Since we use the same set of regressors for the six models (apart from the agent-level "other covs" are different for job-seeker and vacancy estimates), we only display the specifications for h^w and omit the other five equations.

It is important to note that, when using agent-level data from both sides of the same market, the parameters $\{\alpha,\ \beta,\ \alpha_1,\ \beta_1,\ \alpha_1,\ \beta_1\}$ are over-identified because they are estimated twice. This provides a test of our specifications not available to other investigators. The idea is that measurement errors in U and V stocks lead to attenuation in the estimates of the coefficients on U and V, and this causes a gap between the estimated α's and β's from the two sides of the market. For instance, suppose $\log U_{kt}$ is mismeasured. When using job-seeker data, the estimates of $\alpha - 1$ and $\alpha_1 - 1$ in (3.10) and (3.11) are biased towards zero, leading to over-estimates of α and α_1, whereas, when using vacancy data, the attenuation bias leads to underestimates of α and α_1 in (3.13) and (3.14). By symmetrical arguments, the gaps in the estimates in β and β_1 goes the other way.

In the next two sub-sections we explore two important reasons why the stocks of U and V are mismeasured.

3.3.3 Temporal aggregation bias

In keeping with most of the literature, our stocks of job-seekers U_{kt} and vacancies V_{kt} are observed at approximately monthly intervals. The meeting and matching process is continuous, i.e. events such as inflows and outflows occur on daily or even hourly basis, whereas the unemployment and vacancy stocks and outflow of matches are observed over limited frequency. Estimating a continuous process using discrete data leads to a *temporal aggregation bias* (TAB).[2] Based on the assumption that the transition rates between employment and unemployment are constant within the period, two methods have been used to deal with this temporal aggregation bias. Shimer (2012) deals with the temporal aggregation bias by computing the instantaneous transition rates. Coles and Petrongolo (2008) introduce a method to compute the true at-risk stocks corrected for the instantaneous inflows and outflows.

In this subsection, we explain precisely what TAB is, and how Coles and Petrongolo (2008), deal with this bias. They first define, for a given calendar month t, U_t as the beginning-of-month stock, U_{t+1} as the end-of-month stock, u_t as the inflow of job-seekers during the month, and m_t as the outflow, that is, the number of job-seekers who match during the month.[3] These are aggregate measures for a given labour market k.

Coles and Petrongolo obtain a discrete-time representation of the continuous-time matching process by assuming that the inflows of new agents, u_t and v_t are constant within a month. Then total matches during month t are

$$M_t = \int_0^1 h^w U_t e^{-h^w z} dz + \int_{x=0}^1 u_t dx \left[\int_{z=x}^1 h^w e^{-h^w (z-x)} dz \right],$$

[2]This is distinct from *Simultaneity Bias*, another issue investigators in this literature have dealt with. It occurs when the average of the beginning-of-period and end-of-period stocks of U and V are used to capture the changes in the stocks over the month. This is problematic because as the outflow of matches depletes the average stocks, a random shock in matches reduce the stocks, leading to a negative bias in the estimated arguments of the matching function. Then an IV estimator is needed (e.g. Berman, 1997), unlike here – recall that our U_{kt} and V_{kt} are measured at the beginning of each month.

[3]We ignore other outflows in this subsection and we suppress the subscript k throughout.

with a similar expression for vacancies. h^w is the matching rate of the representative job-seeker and h^e is the matching rate of the representative vacancy.[4] The first term represents job-seekers who match from the stock and the second term represents job-seekers who match from the inflow. Evaluating the integrals

$$M_t = h^w U_t \left[\frac{1-e^{-h^w}}{h^w}\right] + h^w u_t \left[\frac{e^{-h^w}-1+h^w}{h^w h^w}\right] \equiv h^w \bar{U}_t \equiv \phi_1 U_t + \phi_2 u_t \tag{3.16}$$

$$= h^e V_t \left[\frac{1-e^{-h^e}}{h^e}\right] + h^e v_t \left[\frac{e^{-h^e}-1+h^e}{h^e h^e}\right] \equiv h^e \bar{V}_t \equiv \phi_3 V_t + \phi_4 v_t, \tag{3.17}$$

with $\{\phi_1, \ldots, \phi_4\}$ defined accordingly. This shows that the number of matches in a calendar month depends on both the beginning-of-period stock U_t and the inflow u_t, and the same for V_t and v_t. The relative importance of the stock to flow depends on the underlying constant hazards h^w and h^e.[5]

We now use these "at-risk" measures of $\bar{U}_t$ and $\bar{V}_t$ to estimate our job-seeker matching model. The *correct* specification replaces $\{U_t, V_t\}$ by $\{\bar{U}_t, \bar{V}_t\}$ in Equation (3.10), which becomes:

$$\log h^w_{in} = \kappa + (\alpha - 1)\log \bar{U}_t + \beta \log \bar{V}_t + \log \tau_{in} + \theta_n + \text{other covs} + \varepsilon_i. \tag{3.18}$$

Using Equations (3.16) and (3.17) above,

$$\log h^w_{in} = \kappa + (\alpha - 1)\log(\phi_1 U_t + \phi_2 u_t) + \beta \log(\phi_3 V_t + \phi_4 v_t) + \log \tau_{in} + \theta_n + \text{other covs} + \varepsilon_i. \tag{3.19}$$

[4] In this subsection, the hazards h^w and h^e are re-scaled to monthly rates.

[5] Coles and Petrongolo provide some intuition: letting $h^w \to 0$ implies $\bar{U} \to U_t + \frac{1}{2}u_t$, i.e. given that the matching hazard is zero, each newly unemployed job-seeker in month t is, on average, unemployed for exactly half of the month, assuming the uniform entry rate. In contrast, letting $h^w \to \infty$ implies each job-seeker is only unemployed for an instant, and $h^w \bar{U} \to U_t + u_t$, i.e. given that all job-seekers immediately find work, total matches equal the total number of job-seekers who search at any point in the month.

A linear approximation leads to the log-linear model:

$$\log h_{in}^{w} = \kappa + \pi_1 \log U_t + \pi_2 \log u_t + \pi_3 \log V_t + \pi_4 \log v_t + \log \tau_{in} + \theta_n + \text{other covs} + \varepsilon_i. \tag{3.20}$$

The relationship between the π's and the model parameters is found by differentiating Equation (3.19):

$$\frac{\partial \log h_{in}^{w}}{\partial U_t / U_t} = \frac{(\alpha - 1)\phi_1 U_t}{\phi_1 U_t + \phi_2 u_t} = \pi_1, \qquad \frac{\partial \log h_{in}^{w}}{\partial u_t / u_t} = \frac{(\alpha - 1)\phi_2 u_t}{\phi_1 U_t + \phi_2 u_t} = \pi_2. \tag{3.21}$$

From this, $\alpha - 1 = \pi_1 + \pi_2$. Similarly, $\beta = \pi_3 + \pi_4$. In other words, Equation (3.20) shows how we estimate α and β unbiasedly: we simply add together the estimates on U_t and u_t to get $\alpha - 1$ and add together the estimates on V_t and v_t to get β.[6] This deals with the obvious reverse causality whereby the at-risk stocks depend on the hazard rates. The intuition is as follows. Equation (3.16) tells that h^w determines the relative importance of U_t and u_t in determining the number of matches, via π_1 and π_2 respectively; however, h^w has no effect on the sum $\pi_1 + \pi_2$.

Equation (3.19) also explains why TAB occurs. TAB is the omitted-variables bias that comes about because the two inflow terms u_t and v_t have been ignored. This makes perfect sense in that these inflows lead to some unmodelled matches during the subsequent month. Indeed, under some simplifying assumptions, we can derive the following expression for the bias in $\widehat{\alpha - 1}$:

$$\text{Bias}(\widehat{\alpha - 1}) = -(\alpha - 1)\left[\frac{\phi_2 u_t}{\phi_1 U_t + \phi_2 u_t}\right]\left[1 - \frac{\text{Cov}(\log U_t, \log u_t)}{\text{Var}(\log U_t)}\right].$$

The first fraction lies in $[0, 1]$, the second in $[-1, 1]$, and so the closer each is to zero, the smaller the bias. The first gets smaller the shorter the aggregation period (here, a month);

[6] In Appendix 3.A we show that this method produces unbiased estimates using simulated estimations of the aggregate contact function.

the second is a characteristic of the data. Thus the bias in $\widehat{\alpha - 1}$—itself negative—is attenuated towards zero, from above because of the sign change. Symmetrical considerations apply to Bias$(\widehat{\beta})$.

To conclude this subsection, we deal with TAB by adding the flows $\log u_{kt}$ and $\log v_{kt}$ as extra covariates to Equation (3.10). The same applies to the other five equations in the base model above. In other words, adding flows $\log u_{kt}$ and $\log v_{kt}$ to Equations (3.10)–(3.15), corrects the TAB.

3.3.4 Geographic spill-overs

In this subsection we discuss our empirical approach to account for the cross-market search. For job-seekers, meeting and matching with a vacancy in the local labour market is easier and less costly than with a distant vacancy. However a job-seeker is not constrained by the border of their local labour market; rather, an increase in the stock of vacancies in a neighbouring labour market can also increase the matching hazard, although to a less extent than a same increase in the stock of local vacancies. The fact that neighbouring labour markets are partially integrated makes it important to adopt a flexible framework which includes distant stocks, yet allows the effects of the local stocks and stocks in neighbouring markets to differ.

Our measures of job-seeker and vacancy stocks are defined at the level of local labour market k. Denote by r the region, which is the set of neighbouring markets that includes market k, we assume that spill-overs occur between the markets within the same region (more details are given in Section 3.4.4), yet agents are not completely indifferent between the local market and neighbouring markets (in which case the correct specification is at region level).

To account for this spillover, we specify the effective job-seeker stock for agents in market k as

$$\tilde{U}_k = U_k + \delta^w U_{r(!k)},$$

where $U_{r(!k)} \equiv U_r - U_k$ is the U stock in other markets of region r, which is discounted by $\delta^w \in (0,1]$, and the vacancy stock as

$$\tilde{V}_k = V_k + \delta^e V_{r(!k)},$$

where $V_{r(!k)} \equiv V_r - V_k$ is V stock in other markets of region r, which is discounted by $\delta^e \in (0,1]$ (similar specifications are used in Burda and Profit, 1996; Burgess and Profit, 2001; Petrongolo and Manning, 2017). Notice that δ's measure the extent of the spill-over and are extra parameters to be estimated from data. For instance, from the job-seeker's perspective, δ^w can be interpreted as a measure of the conversion rate for the congestion effect caused by a job-seeker in a neighbouring market relative to that caused by a job-seeker in the home market. When $\delta^w = 0$ there are no spill-overs, so U_{kt} is the correct specification, when $\delta^w = 1$ the region acts as a single labour market so U_{rt} is the correct specification, and when $0 < \delta^w < 1$, there is measurement error if either U_{rt} or U_{kt} are used.

To estimate our models, we approximate the stocks $\log \tilde{U}_k$ and $\log \tilde{V}_k$ by[7]

$$\log \tilde{U}_k = \log(U_k + \delta^w U_{r(!k)}) \approx \log U_r + \frac{1-\delta^w}{\delta^w}\frac{U_k}{U_r}, \tag{3.22}$$

$$\log \tilde{V}_k = \log(V_k + \delta^e V_{r(!k)}) \approx \log V_r + \frac{1-\delta^e}{\delta^e}\frac{V_k}{V_r}. \tag{3.23}$$

Plugging into Equation (3.10) we get the specification which account for the spill-overs,

$$\log h^w_{in} = \kappa + \phi_1 \log U_{rt} + \phi_2 \frac{U_{kt}}{U_{rt}} + \phi_3 \log V_{rt} + \phi_4 \frac{V_{kt}}{V_{rt}} + \log \tau_{in} + \theta_n + \varepsilon_i + \text{other covs.} \tag{3.24}$$

To get the TAB corrected version of this specification, the same considerations apply to the inflows, so the effective inflows for market k, $\log \tilde{u}_k$ and $\log \tilde{u}_k$, are given by

$$\log \tilde{u}_k = \log(u_k + \delta^w u_{r(!k)}) \approx \log u_r + \frac{1-\delta^w}{\delta^w}\frac{u_k}{u_r}, \tag{3.25}$$

$$\log \tilde{v}_k = \log(v_k + \delta^e v_{r(!k)}) \approx \log v_r + \frac{1-\delta^e}{\delta^e}\frac{v_k}{v_r}. \tag{3.26}$$

Replacing $\{\log U_{kt}, \log V_{kt}, \log u_{kt}, \log v_{kt}\}$ with linearized $\{\log \tilde{U}_{kt}, \log \tilde{V}_{kt}, \log \tilde{u}_{kt}, \log \tilde{v}_{kt}\}$, the model corrected for TAB, Equation (3.20), becomes

$$\begin{aligned}\log h^w_{in} =& \kappa + \pi_{11} \log U_{rt} + \pi_{12} \frac{U_{kt}}{U_{rt}} + \pi_{21} \log u_{rt} + \pi_{22} \frac{u_{kt}}{u_{rt}} \\ &+ \pi_{13} \log V_{rt} + \pi_{14} \frac{V_{kt}}{V_{rt}} + \pi_{23} \log v_{rt} + \pi_{24} \frac{v_{kt}}{v_{rt}} + \log \tau_{in} + \theta_n + \text{other covs} + \varepsilon_i.\end{aligned} \tag{3.27}$$

[7] To get this approximation,

$$\log(U_k + \delta^w U_{r(!k)}) = \log[\delta^w (U_k + U_{r(!k)}) + (1-\delta^w)U_k] = \log\left[(U_k + U_{r(!k)})(\delta^w + \frac{(1-\delta^w)U_k}{U_k + U_{r(!k)}})\right]$$

$$= \log U_r + \log \delta^w + \log(1 + \frac{1-\delta^w}{\delta^w}\frac{U_k}{U_r}) \approx \log U_r + \frac{1-\delta^w}{\delta^w}\frac{U_k}{U_r} + \log \delta^w,$$

i.e. the approximation works if $\frac{1-\delta^w}{\delta^w}\frac{U_k}{U_r}$ is close to zero ($\log(1+x) \approx x$); $\log \delta$'s are omitted from the expressions as they are constant.

The relationships between π's and the model parameters are

$$\pi_{11}+\pi_{21}\equiv(\alpha-1),\quad \pi_{12}+\pi_{22}\equiv(\alpha-1)\frac{1-\delta^w}{\delta^w},\quad \pi_{13}+\pi_{23}\equiv\beta,\quad \pi_{14}+\pi_{24}\equiv\beta\frac{1-\delta^e}{\delta^e}.$$

To conclude the section, Equation (3.27) and its vacancy counterparts, and the analogous specifications for contact hazards and matching probabilities are our complete specifications, which account for TAB and geographic spill-overs.

3.4 Data

The data are from a Matching Service computer system for the county of Lancashire in the UK for March 1988 to June 1992.[8] The system records information on all young job-seekers aged 15-18 (i.e. school-leavers) in the County, vacancies posted by employers, job interviews and matches between young people and those vacancies.[9] The primary purpose of the system is to facilitate search from both sides of the youth labour market: job-seekers can search for vacancies using specified search criteria such as occupational ambitions, preferred locations and qualifications, and the Matching Service staff can search for potential candidates satisfying the requirements of the employers. The data record whether the interview takes place, which is synonymous for a contact, and whether it results in a hire synonymous for match.

[8]Lancashire is in North West England, and excludes large metropolitan cities (Liverpool and Manchester). The Matching Service is actually called Lancashire Careers Service, is funded by government, and its responsibilities are to provide free employment service and vocational guidance to employers and young people. These data are very similar to those analyzed by Berman (1997) for Israel, yet the difference is that their data and analysis are at market aggregate level.

[9]This Matching Service data has been used in previous empirical work to estimate the determinants of matches conditional on contacts (Andrews et al., 2001), the duration of employer search (Andrews et al., 2008) and to test whether the matching rate is consistent with the predictions of stock–flow matching (Andrews et al., 2013).

Each job-seeker using the Matching Service is uniquely identified.[10] Information on job-seekers (such as their current employment status, occupational preferences, qualifications, interviews, etc.) was recorded on a live database from which snapshots were taken approximately each month.[11] Joining together these monthly snapshots enables us to construct a sequence of labour market spells, job interviews and the outcomes of those interviews. We restrict our attention to the search activities of unemployed job-seekers. Those who are on training programmes or in education are excluded from the analysis because, for these job-seekers, the exact date on which search started is less clear than for those who start a spell of unemployment.

The data also contain a record of all job vacancies posted to the Matching Service. Information on vacancies includes occupation, responsibilities and employer information selection criteria and the date at which vacancies open and close.[12] Most vacancies are single positions, but some are multiple positions within the same employer (see Andrews et al., 2008, for a fuller explanation). The proportion of job-seekers who match to vacancies posted via the Matching Service is approximately 25% of all job matches.[13]

Job-seeker interviews and outcomes can be linked to these vacancies, giving us a unique two-sided view of the search process. The interviews and hires are uniquely identified by the pair of agents, since we do not observe any repeated interviews or hires

[10] Strictly speaking, they are identified only within their local office. A job-seeker who moves from one office to another will appear as a new record.

[11] Between March 1988 and June 1992 we have 45 snapshots rather than 51. The median number of days between snapshots is 32.

[12] Most of the vacancies are entry-level, low-skill vacancies, and the top five most frequent job titles are "office junior", "sales assistant", "general assistant", "junior clerk", and "apprentice hairdresser". We also observe wage but only for a minority of the vacancies (44%), so this information cannot be used. The rest of vacancies either do not advertise the wage or have negotiable wages.

[13] However, since employers can choose whether or not to use the Matching Service as a recruitment channel, an important sample issue is the extent to which our sample of vacancies is representative of all vacancies which posted in this labour market. In Table 3.C2 in Appendix 3.C we summarise the characteristics by how spells end. We show that among the job-seekers who end their search spell with a job, those who match with job vacancies via the Matching Service have similar characteristics as those who match outside the the service, and that among the vacancies which are filled, those match with job-seekers via the Matching Service have similar characteristics as those filled by others.

(e.g. temporary layoffs) in our data. Figure 3.1 illustrates the organisation of the data. Our aim is to model search duration, which is the time from the start of the spell to the final successful interview. For one job-seeker 1 who contacts two vacancies, the spell of search starts on date t_3 in month 1. In month 2, job-seeker 1 has an interview with vacancy 2 (c_{12}) and is unsuccessful; she has a second interview with vacancy 1 (c_{11}) which results in a match (m_{11}). She then starts work on date t_7. The date of the final interview is not known precisely, yet we know it is between t_4 and t_6 when the snapshots are taken, and before t_7 – we approximate it using t_6 since $t_6 < t_7$ (in case $t_7 < t_6$ we use t_7 for higher precision), i.e. the job-seeker 1's search spell ends on t_6.

On the other side of the market, vacancy 1 opens on date t_5 and has three interviews with job-seekers, of which the final one is with job-seeker 1 and results in the hire so its search spell also ends on t_6. Vacancy 2 opens on date t_2, has two interviews in month 2 (the first of which is with job-seeker 1 and neither of them results in a hire). This vacancy does not match with a job-seeker within the system, so is censored on date t_8.[14]

3.4.1 Creating estimation panels

As explained in Section 3.3.2, we create two panel datasets, one whose unit of observation is a job-seeker elapsed month and the other at the vacancy elapsed-month level, and we estimate hazards and matching probabilities using standard count-data techniques. This is because we do not observe exactly when each interview took place, only the month in which it occurred. However, we do observe the start and end date of each spell dated to the nearest day,[15] and use this information to calculate more precise measures

[14]This occurs if the vacancy matches with a job-seeker outside the system (e.g. if the job-seeker makes a direct approach to the vacancy) or if the vacancy is withdrawn from the market by the employer.

[15] The only inaccuracy here comes about in the final month of search, because there may be a gap between the date of the successful contact and the date on which the snapshot was taken (e.g. for job-seeker 1 and vacancy 2 in Figure 3.1, this is the gap between the contact (c_{11}) and the snapshot date t_5).

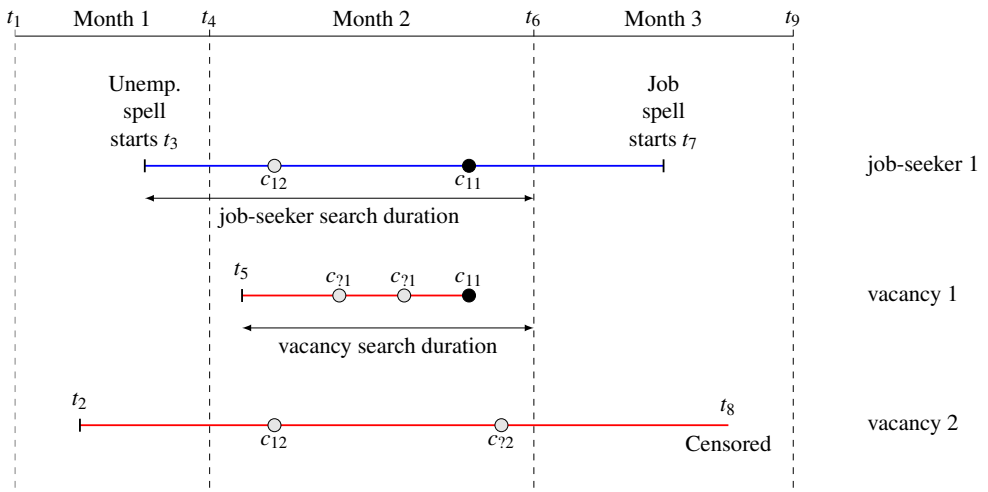


Fig. 3.1 Organisation of the Matching Service data. The search spells of job-seeker 1 and vacancies 1 and 2 are displayed by the blue line and the red lines. An interview between job-seeker i and vacancy j is denoted by c_{ij} (agents who are not displayed are not labeled and hence are represented by question marks). Grey circles indicate unsuccessful interviews, black circles represent successful interviews which result in a hire. Dates which are known are labeled chronologically t_1, t_2, etc. "Months" are not of equal lengths because the date on which snapshots were taken varied across months.

of time-at-risk days within a month. The precise spell duration is also used to measure job-seeker and vacancy stocks and inflows, as we explain later in Section 3.4.3.

The data recorded in our job-seeker and vacancy panels are the counts of contacts and matches, the time spent searching in each elapsed month (approximately 30 days except for the first and last months of the search spell), and the market and agent characteristics ($\{C_{in}, M_{in}, \tau_{in}, x'_{in}\}$ and $\{C_{jn}, M_{jn}, \tau_{jn}, x'_{jn}\}$ respectively). For example, for job-seeker 1 illustrated in Figure 3.1, the data would be recorded as $\{0,\ 0,\ t_4 - t_3,\ \ldots\}$ in month 1 (i.e. no contact and no match over the $(t_4 - t_3)$ days of search), and $\{2,\ 1,\ t_6 - t_4,\ \ldots\}$ in month 2 (i.e. two contacts and one match over the $(t_6 - t_4)$ days of search), for the two months she is observed. The data for vacancy 1 would be recorded as $\{3,\ 1, t_6 - t_5,\ \cdots\}$ in month 2 for the one month it is observed (i.e. three contacts and one match over the $(t_6 - t_5)$ days of search).

The job-seeker panel and the vacancy panel share the same pool of contacts and matches, because they describe two sides of one well-defined labour market. We focus on contacts and matches made between unemployed job-seekers and job vacancies. We recognise that job-seekers may also contact and match with vacancies which are not posted on the system, or with training vacancies. And vacancies may contact and match with job-seekers who are not unemployed (such as job-seekers who are in education) or who do not use the Matching Service system. These spells are censored (e.g. vacancy 2 in Figure 3.1), but are included in the risk set until that point.

3.4.2 Sample statistics

The job-seeker panel contains 38,223 unemployed job-seekers who have 49,093 search spells and 166,625 search months. These unemployed job-seekers make 36,451 contacts to job vacancies registered on the system, which result in 3,095 matches. The vacancy data panel contains 12,550 vacancy orders[16] with 17,510 vacancy spells and 26,764 search months. The average durations of job-seeker and vacancy search spells are 85 and 73 days, and the median durations are 55 and 23 days, respectively. The average observable characteristics of the spells are summarised in Table 3.1.

The average number of total contacts for job-seeker spells is 0.74. Less than 20% of the job-seeker months have contacts, which reflects the fact that the market condition is in general very slack. The average number of contacts for a vacancy spell is 2.08. Among all job-seekers, 6% match within the system, 21% match with a vacancy outside the system, and 73% of job-seeker spells are censored (the job-seeker becomes inactive). Among all vacancies, 18% match within the system, 42% match with a job-seeker outside the system, and 40% of vacancy spells are censored (the vacancy is withdrawn). The average observable characteristics of job-seekers and vacancies who match within

[16] A vacancy order may contain several vacancy positions with identical characteristics, e.g. an employer may advertise for "three apprentice mechanics".

the system is very similar to those who match outside the system (see Table 3.C2 in Appendix 3.C).

Table 3.1 Summary statistics of job-seeker and vacancy panels

Variable	Mean (Median)	Std. Dev.
(a) Job-seeker panel		
Duration in months	3.39 (3)	2.85
Duration in days	85 (55)	93.75
Total contacts	0.74	1.56
Prop. matched through LCS	0.06	
Prop. Male	0.58	
Prop. Non-white	0.06	
Prop. Receiving subsidy	0.07	
Prop. with higher-level quals[a]	0.14	
Number of job-seekers	38,223	
Number of unemployment spells	49,093	
Number of job-seeker months	166,625	
(b) Vacancy panel		
Duration in months	2.14 (1)	2.27
Duration in days	73.02 (23)	119.48
Total contacts	2.08	4.73
Prop. matched through LCS	0.18	
Prop. skilled	0.54	
Prop. non-manual	0.54	
Prop. Written application	0.27	
Prop. Large firms ($>$ 50 employees)	0.22	
Number of employers	4,999	
Number of vacancy orders[b]	12,550	
Number of individual vacancies	17,510	
Number of vacancy-order months	26,764	

[a] With more than three O-level qualifications.
[b] 13% of the vacancy orders placed by the employers contain multiple positions with identical characteristics.

3.4.3 Stocks and inflows

The stocks of job-seekers and vacancies at the beginning of each month and inflows over the subsequent month are generated by aggregating information of search duration of individual agents. Specifically, the beginning-of-month U stock of month t and $t+1$ are

linked by

$$U_{t+1} = U_t + u_t - u_t^- = U_t + u_t - M_t - u_t^o,$$

where the count of job-seekers who enter the market during month t is the total monthly inflow u_t, the count of those exit the market during month t is the total monthly outflow u_t^-. The monthly outflow u_t^- includes u_t^o, the total monthly censored outflow, and M_t, the total number of matches during the month. M_t appears in the completely analogous expression for the vacancy stock:[17]

$$V_{t+1} = V_t + v_t - v_t^- = V_t + v_t - M_t - v_t^o,$$

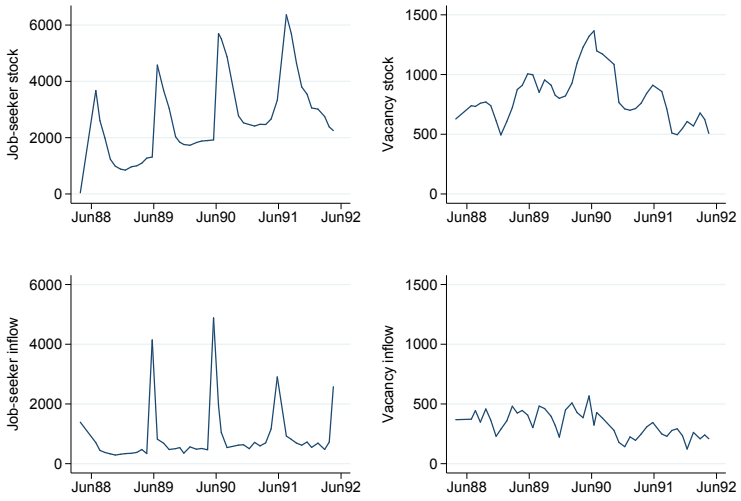
where v_t is vacancy inflow, v_t^- is the vacancy outflow, among which v_t^o are the censored outflow.

In Figure 3.2a we plot the aggregate time-series for $\{U_t, V_t, u_t, v_t\}$. There are large peaks in unemployed stocks arising from young people leaving school between May and August each year, which, of course, is when employers post their vacancies. There is a similar annual variation in the data for vacancy stocks, but less pronounced. This so-called recruitment cycle is important for identification, because it generates a lot of time-series variation in the data. In Figure 3.2b, we plot the aggregate time-series for our three dependent variables, namely contacts, C_t, matches M_t and the matching probability M_t/C_t. It is clear that the aggregate number of matches varies positively with the stocks of U_t and V_t, as it should, mimicking the movement of V_t more closely than U_t, again as it should in a slack labour market. The matching probability exhibits less variation through the sample period.

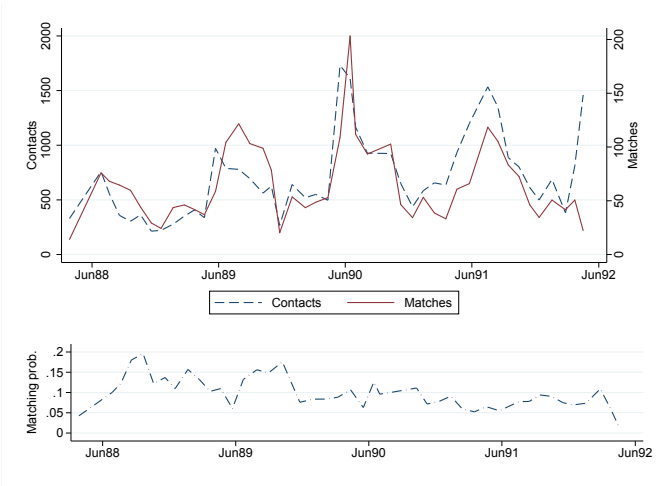
[17] This is only true when we aggregate over the entire Lancashire. When considering the local labour market k, due to cross market search, the total number of matches which occur to the job-seekers located in k is different from the total number of matches occur to vacancies located in k. This is illustrated the next subsection.

Fig. 3.2 Aggregate monthly stocks, inflows, contacts and matches

(a) Job seeker and vacancy stocks $\{U_t, V_t\}$ and inflows $\{u_t, v_t\}$



(b) Matches M_t, contacts C_t, and matching probability μ_t



3.4.4 Defining local districts and regions

To estimate our models with geographic spill-overs, we need to define our local labour markets and regions. Given that we have agent-level search duration and location information, we can construct the measures of the stocks and inflows of job-seekers and vacancies at different levels of spatial aggregation. The most aggregated level is the whole of Lancashire. All of the meetings and matches we observe are by definition between the aggregate stocks of job-seekers and vacancies if we treat Lancashire as one aggregate market. However, the size and geographic features Lancashire imply that it is unlikely that the opening of a vacancy anywhere in Lancashire matters equally for a job-seeker in a particular location. And at the least aggregated level, Lancashire comprises of 14 distinct districts which are towns/cities and their surrounding countryside.[18]

The operation of the Matching Service makes each of the 14 districts a natural choice for our local labour market k: the system has one or two offices within each area, of which the staff meet with the local job-seekers and employers, and might prioritise the arrangement of interviews for local agents registered with the office compared to those from other offices. In addition, the transportation costs are lower, as buses run more frequently within the district.

The towns can also be grouped into three larger regions: West, Central and East. The traveling times between towns in each group are typically within one hour and by direct buses, which makes commuting a feasible option for the young unemployed. Therefore we use this grouping to define the regions r, within which the spill-overs take place. See Figure 3.B1 in Appendix 3.B for a map of Lancashire, which displays the districts and regions. We plot the stocks and inflows for the 14 local markets $\{U_{kt},\ V_{kt}, u_{kt},\ v_{kt}\}$ and the three regions $\{U_{rt},\ V_{rt}, u_{rt},\ v_{rt}\}$ in Figure 3.C1 in Appendix 3.C.

[18] Strictly speaking it is still possible to further disaggregate into smaller areas because the data records the postcodes of the job-seekers and firms; we do not exploit more detailed location information for the purpose of this study.

We summarise the amount of spill-overs at the aggregate level. Amongst all contacts, 69.8% occur within the same local labour market (i.e. the i, j pair both locate in the same k), whereas 92.7% occur within the the same region (the i, j pair both locate in r). The corresponding figures for matches are 79.2% and 95.7% respectively. Clearly, there are substantial spill-overs between the districts within the region, with job-seekers residing in one town and contacting employers in another. Also, there are only negligible spill-overs between the regions, so is unnecessary to incorporate the stocks at more aggregated level.

The evidence of spill-overs justifies our specification given in Section 3.3.4. Measurement errors will arise if we use the local market stocks and inflows or if we use region stocks and inflows. The stocks in the neighboring markets are likely to have a significant yet weaker effect on the hazards compared to the stocks in the local market. An average job-seeker clearly prefers a local vacancy to a distant one: although the size of local market stock is merely roughly a quarter of the region stock, nearly 70% of contacts are with local agents.

3.5 Results

3.5.1 Comparison of specifications

In Table 3.2 we report estimates of α and β from Equation (3.10) (and (3.13)) in column (1), i.e. the simplest specification which uses as covariates the stock of job-seekers and vacancies at the beginning of each month. As explained, we expect that this specification will suffer from measurement error which will manifest itself in (i) overestimating α from the job-seeker data and underestimating α from the vacancy data; and (ii) underestimating β from the job-seeker data and overestimating β from the vacancy data. Comparing panels (a) and (b) of Table 3.2 confirms this prediction. The estimate of

α from the job-seeker matching hazard h^w ($\widehat{\alpha_w}$) is twice as large as the estimate of α from the vacancy matching hazard h^e ($\widehat{\alpha_e}$) ($\widehat{\alpha_w} = 0.944$ compared to $\widehat{\alpha_e} = 0.483$, and the absolute gap $\mid \widehat{\alpha_w} - \widehat{\alpha_e} \mid$ is 0.461). Estimates of β from the vacancy matching hazard are five times larger than estimates from the job-seeker matching hazard ($\widehat{\beta_e} = 0.986$ compared to $\widehat{\beta_w} = 0.154$, and the absolute gap $\mid \widehat{\beta_w} - \widehat{\beta_e} \mid$ is 0.832).[19]

In column (2) and (3) we report the estimates correcting for TAB but not spillovers (Equation 3.20 in Section 3.3.3), and the ones correcting for spillovers but not TAB (Equation 3.24 in Section 3.3.4). Both specifications brings $\hat{\alpha}$ and $\hat{\beta}$ from the job-seeker and vacancy sides closer together compared to column (1): for α the gaps are $\mid \widehat{\alpha_w} - \widehat{\alpha_e} \mid = \mid 0.448 - 0.449 \mid = 0.001$ in column (2) and $\mid 0.933 - 0.498 \mid = 0.435$ in column (3), both smaller than 0.483 in column (1); for β the gaps are $\mid \widehat{\beta_w} - \widehat{\beta_e} \mid = \mid 0.521 - 0.781 \mid = 0.260$ in column (2) and $\mid 0.349 - 1.016 \mid = 0.667$ in column (3), both smaller than 0.832 in column (1).

In column (4), we estimate the specification which corrects for both the TAB and spill-overs, as given by Equation (3.27). The raw coefficients and the estimates of parameters for the degree of spill-overs $\{\delta^w, \delta^e\}$ are displayed in Table 3.D1 in Appendix 3.D. Adding inflows as covariates and accounting for the stocks in the neighbouring markets pushes the estimates of α and β much closer together, as predicted. We now have $\hat{\alpha} = 0.471$ from the job-seeker hazard and $\hat{\alpha} = 0.488$ from the vacancy hazard; $\hat{\beta} = 0.554$ from the job-seeker hazard and $\hat{\beta} = 0.700$ from the vacancy hazard. The estimates from the two sides of the market are much more similar (with gaps of $\mid \widehat{\alpha_w} - \widehat{\alpha_e} \mid = 0.017$ and $\mid \widehat{\beta_w} - \widehat{\beta_e} \mid = 0.146$) using our fully corrected specification (column (4)) than using any other specifications (column (1)–(3)).

[19] In fact this discrepancy has been documented in other studies using market-level estimates, e.g. Burgess and Profit (2001) also report very different estimates of α and β using the unemployment outflow and vacancy filling rate as the dependent variables.

Table 3.2 Estimates of α and β from job-seeker and vacancy panels

(a) Job-seekers h^w

	(1)	(2)	(3)	(4)
α	0.944 (0.066)	0.448 (0.090)	0.933 (0.069)	0.471 (0.095)
β	0.154 (0.076)	0.521 (0.091)	0.349 (0.112)	0.554 (0.122)
$\alpha+\beta$	1.098 (0.093)	0.969 (0.114)	1.281 (0.122)	1.025 (0.139)
TAB corrected	No	Yes	No	Yes
Spill-overs corrected	No	No	Yes	Yes
Log Likelihood	-16,082	-16,036	-16,077	-16,041
Obs.	166,625	166,625	166,625	166,625

(b) Vacancies h^e

	(1)	(2)	(3)	(4)
α	0.483 (0.080)	0.449 (0.106)	0.498 (0.084)	0.488 (0.112)
β	0.986 (0.090)	0.781 (0.108)	1.016 (0.136)	0.700 (0.147)
$\alpha+\beta$	1.469 (0.109)	1.230 (0.135)	1.514 (0.141)	1.189 (0.167)
TAB corrected	No	Yes	No	Yes
Spill-overs corrected	No	No	Yes	Yes
Log Likelihood	-9,396	-9,388	-9,394	-9,382
Obs.	26,764	26,764	26,764	26,764

* All regressions include elapsed duration dummies, dummies for local labour market k, four quarter dummies, five year dummies, and five year dummies interacted with local labour market k, "own" agent-level covariates and market-level covariates for both sides of the market. See Section 3.3 for fuller details.

+ All models have standard errors clustered at unemployment and vacancy spell level.

3.5.2 Decomposition of the matching function

Our corrected specification passes the over-identification check, so it allows us to decompose the effect of labour market tightness into a contact effect λ and a matching probability effect μ. The estimates of λ and μ are displayed in Table 3.3. The columns labeled "Uncorrected" are the estimates without correcting for TAB and spill-overs (as column (1) in Table 3.2), and the columns labeled "Corrected" are corrected for both problems (as column (4) in Table 3.2). The "Uncorrected" models show that, the corrections are also necessary for λ estimates to deliver consistent results from the two sides of the market: $\widehat{\alpha_{1w}} = 1.046$ from job-seeker side, which is over five times larger than $\widehat{\alpha_{1e}} = 0.154$ from the vacancy side, and $\widehat{\beta_{1w}} = 0.098$ from the job-seeker side, which much smaller than $\widehat{\beta_{1e}} = 1.183$ from the vacancy side. For $\hat{\alpha}_2$ and $\hat{\beta}_2$ there are also gaps between the job-seeker and vacancy sides.

We use the corrected specification for the decomposition.[20] On the job-seeker side, $\hat{\alpha} = 0.471$ arises because there is a strong positive effect of job-seeker stocks on the contact rate $\hat{\alpha}_1 = 0.673$ but a negative effect on the matching probability $\hat{\alpha}_2 = -0.201$. And $\hat{\beta} = 0.554$ is also roughly equal to the sum of a strong positive effect of vacancy stocks on the contact rate $\hat{\beta}_1 = 0.384$ and a small and insignificant positive effect on the matching probability $\hat{\beta}_2 = 0.087$. The congestion effect $(\alpha - 1)$ appears to be weaker for contacts than for matches $(\hat{\alpha} - 1 < \hat{\alpha}_1 - 1 < 0)$, and the thick-market effect β is also weaker for contacts than for matches $(0 < \hat{\beta}_1 < \hat{\beta})$. In other words, in slack labour markets there are more interviews, but each interview has a lower probability of success.

The estimates on the vacancy side do not decompose quite so neatly: $\hat{\alpha} = 0.488$ is slightly lower than the sum of $\hat{\alpha}_1 = 0.856$ and $\hat{\alpha}_2 = -0.220$, and the estimates of $\hat{\beta} = 0.700$ is slightly higher than the sum of $\hat{\beta}_1 = 0.391$ and $\hat{\beta}_2 = 0.094$. However, the standard errors are also slightly larger for the vacancy estimates.

[20] For λ and μ, the raw coefficients and estimates of δ's are again in Table 3.D1 in Appendix 3.D.

Table 3.3 Estimates of $\{\alpha,\beta\}$, $\{\alpha_1,\beta_1\}$ and $\{\alpha_2,\beta_2\}$ from job-seeker and vacancy panels, and contacts cross-section

(a) Job-seekers	h^w		λ^w			μ^w	
	Uncorrected	Corrected	Uncorrected	Corrected	Corrected with FE	Uncorrected	Corrected
α	0.944	0.471	1.046	0.673	0.888	−0.024	−0.201
	(0.066)	(0.095)	(0.028)	(0.040)	(0.052)	(0.063)	(0.090)
β	0.154	0.554	0.098	0.384	0.474	0.020	0.087
	(0.076)	(0.122)	(0.032)	(0.051)	(0.070)	(0.071)	(0.118)
$\alpha+\beta$	1.098	1.025	1.144	1.057	1.363	−0.004	−0.114
	(0.093)	(0.139)	(0.041)	(0.059)	(0.079)	(0.086)	(0.137)
Log Likelihood	-16,082	-16,041	-89,504	-89,291	-34,577	-9,844	-9,839
Obs.	166,625	166,625	166,625	166,625	65,431	36,451	36,451

(b) Vacancies	h^e		λ^e			μ^e	
	Uncorrected	Corrected	Uncorrected	Corrected	Corrected with FE	Uncorrected	Corrected
α	0.483	0.488	0.154	0.856	0.902	0.259	−0.220
	(0.080)	(0.112)	(0.043)	(0.059)	(0.090)	(0.124)	(0.141)
β	0.986	0.700	1.183	0.391	0.439	−0.041	0.094
	(0.090)	(0.147)	(0.056)	(0.097)	(0.155)	(0.075)	(0.130)
$\alpha+\beta$	1.469	1.189	1.337	1.247	1.341	0.218	−0.126
	(0.109)	(0.167)	(0.065)	(0.099)	(0.159)	(0.145)	(0.184)
Log Likelihood	-9,396	-9,382	-37,677	-37,246	-12,429	-9,829	-9,807
Obs.[a]	26,764	26,764	26,764	26,764	13,953	36,451	36,451

[*] All regressions include elapsed duration dummies, dummies for local labour market k, four quarter dummies, five year dummies, and five year dummies interacted with local labour market k, "own" agent-level covariates and market-level covariates for both sides of the market. See Section 3.3 for fuller details.

[+] All models have standard errors clustered at unemployment and vacancy spell level.

[a] For random effect estimates, the sample sizes are 166,625 job-seeker months from job-seeker data, 26,764 vacancy months from vacancy data, and 36,451 contacts between pairs of job seekers and vacancies from contact data. The fixed effect estimate for job-seekers has a sample size of 65,431 job-seeker months. It excludes 11,054 unemployment spells (and months) because they only last for one month, and further excludes 24,144 unemployment spells (90,140 months) because they have no contacts in any month. The fixed effect estimate for vacancies has 13,953 vacancy months, with 6,558 spells (and months) excluded because they only last for one month, and 1,792 spells (6,253 months) excluded because they have zero contacts in all months.

The results on returns to scale are generally consistent with our prior expectations. Our estimated returns to scale vary from 1.025 (standard error 0.139) to 1.189 (0.167) for matches, from 1.057 (0.059) to 1.247 (0.099) for contacts, and from -0.114 (0.137) to -0.126 (0.184) for the matching probability. We cannot reject the null of constant returns to scale from either dataset for matching hazards. For meeting hazards, we do not reject constant returns to scale for the job-seeker side, although we do reject constant returns to scale for the vacancy side with a two-sided p-value of < 0.001. Also, $\hat{\alpha} + \hat{\beta}$ always tend to be larger from the vacancy data. This might suggest that the specification can still be improved.[21] We also cannot reject zero returns to scale in the matching probability from both job-seeker and vacancy side.

For the contact rate (λ^w and λ^e) there is a third column, labelled "Corrected with FE", which displays the Fixed Effect estimates using the corrected specification, controlling for unobserved heterogeneity. Estimated returns to scale are larger from both sides of the market using the FE estimator, with $\hat{\alpha}_2 \approx 0.9$ and $\hat{\beta}_2$ from 0.4 to 0.5. We reject constant returns to scale in the FE contact rates. The difference between the fixed effect and random effect estimates suggest that there is some predictable correlation between the job-seeker or vacancy fixed effect and the stocks U and V. Notice that the samples for the Fixed Effects estimates tend to be those who are searching for longer. This positive scale effect in the matching and contacting function suggests that for the agents who fail to match immediately it might be easier to find a match in larger markets than in the smaller markets.

[21] For instance, we recognise that there are substantial outflows which are not matches within the Matching Service system and are asymmetric for the job-seeks and vacancies. This means it might be necessary to take the censored outflows into account in the correction of TAB.

3.5.3 Decomposition of hazard rates

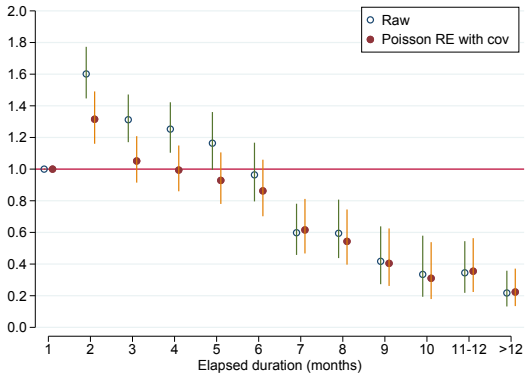
In this section we report estimates of the hazard rates for job-seekers and vacancies and their decomposition into the change in the arrival rate of contacts and the change in the probability of a match conditional on a contact. Each panel of Figure 3.3 compares the raw hazard and estimates of the hazard from a Poisson random effects model which includes a complete set of covariates as shown in (3.27). We normalise the hazard rate in the first month of search to one so that all estimates are reported as incidence rate ratios (Table 3.C1 in Appendix 3.C reports the raw hazard rate). For example, the estimated raw hazard for month 2 in panel (a) shows that the estimated rate for month 2 is 1.6 times of that for month 1 in terms of raw statistics, and 1.2 times of month 1 in terms of the estimates.

We note that the inclusion of covariates and the use of a random effects model makes little difference to the estimated hazards in all six panels. To some extent this is surprising, because if selection on observable characteristics matters then we would expect the hazard rates to decline less steeply after controlling for those characteristics. But this is not the case.

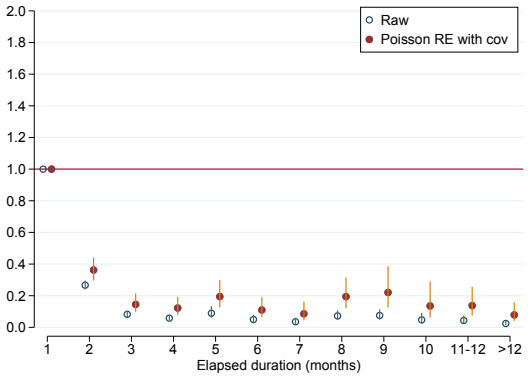
In panel (a) the matching hazard rate for job-seekers jumps after the first month and declines steadily thereafter. In panel (b) the matching hazard rate for vacancies falls very sharply after the first month, falls again after the second month and is quite flat thereafter. Panels (c) and (e) show that the jump in the hazard for job-seekers occurs because the matching probability jumps between months 1 and 2,[22] and that the subsequent fall is entirely because of a declining contact rate. Panels (d) and (f) show that the decline in the vacancy hazard is entirely because of a declining contact rate.

[22]For job-seekers the matching probability is low in the first month of search compared to later months, which might reflect the process that the job-seekers adapt to realistic expectations about the market condition. In addition, as noted in Andrews et al. (2001), this can also result from an 8-week benefit payments received by 16- and 17-year-olds which makes them more selective in the first month.

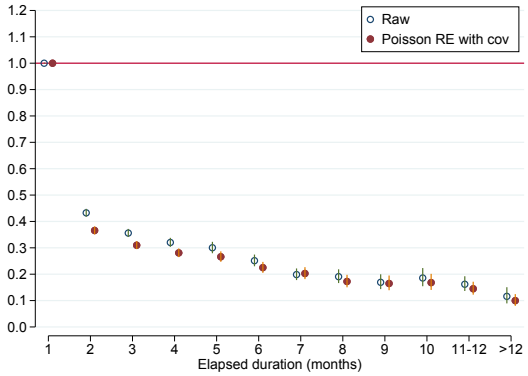
Fig. 3.3 Comparison of raw and conditional hazard rates



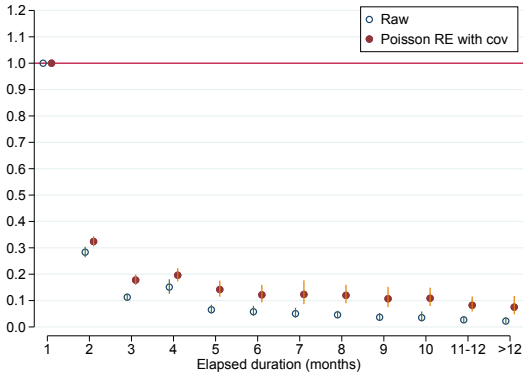
(a) Job-seeker matching hazard



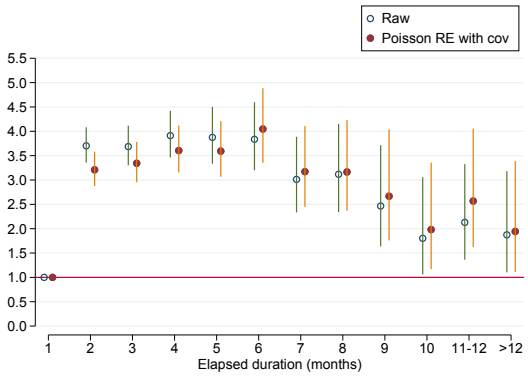
(b) Vacancy matching hazard



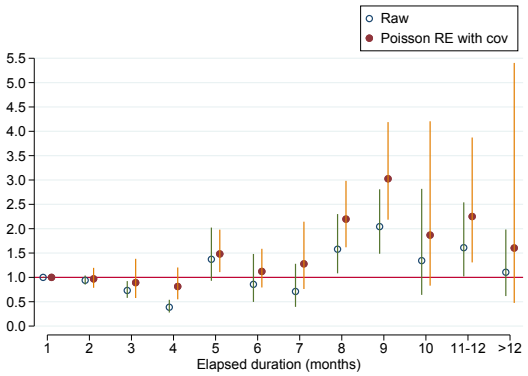
(c) Job-seeker contact hazard



(d) Vacancy contact hazard



(e) Job-seeker matching probability



(f) Vacancy matching probability

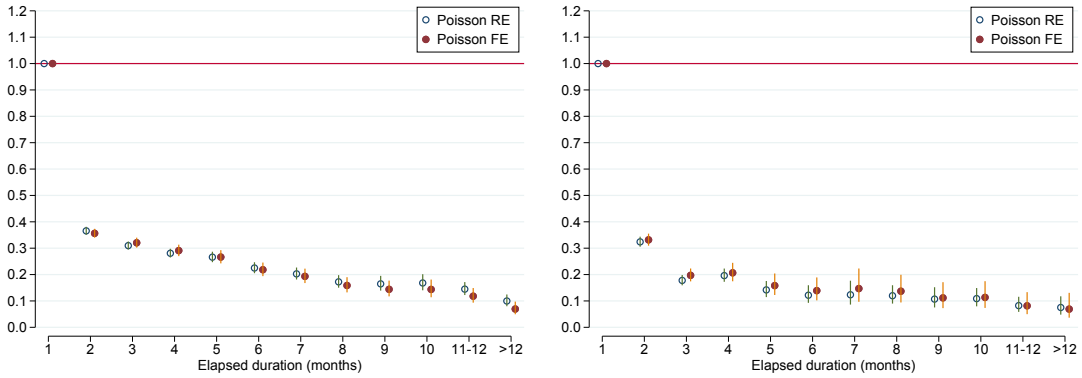
There are three potential explanations for the fall in the contact rate observed on both sides of the market. The first explanation is the between-agent quality change as a result of sample selection. The second explanation is the within-agent quality change due to changes in characteristics (e.g. loss of human capital). And the third explanation is some non-random matching mechanism.

We are able to rule out the first explanation by providing a genuine negative duration dependence, exploiting the repeated observations on contacts for the same job-seekers and vacancies. This allows us to estimate fixed-effect Poisson models. Results are shown in Figure 3.4. Again surprisingly, the estimates are very similar for the fixed effect and random effect estimates, which suggests that the unobserved heterogeneity does not explain the duration patterns of contact hazards.

The finding is consistent with the stock-flow matching model. If there is a marketplace in which job-seekers and vacancies can contact each other quickly, we would expect very high initial contact rates. But once this initial period is over, agents only contact new entrants on the other side of the market, and the contact rate therefore falls sharply. Since the data is indeed from one market in which the identities of agents on both sides of the market is observable to the matching service, this kind of stock-flow effect is plausible. We further explore this explanation in more detail in Chapter 4.

Of course, these contact hazards are also consistent with the random matching model if one assumes there is genuine duration dependence as a result of the "quality" of job-seekers and vacancies declining with duration. The latter seems unlikely, however, over such a short period, as the largest drop in contact rate occur in the second month. Also, we do not think that vacancies' characteristics decay with duration, although very similar patterns are present in the vacancy contact hazards.

Fig. 3.4 Baseline hazards for meetings, conditional on covariates and fixed-effects. The inclusion of FE does not change the underlying shape of the meeting hazard



(a) Job-seeker contact hazard δ^u

(b) Vacancy contact hazard δ^v

3.5.4 Decomposition of the effect of characteristics

In Table 3.4 we report estimates of the effect of job-seeker and vacancy characteristics on the matching rate (h^w and h^e) and its component parts, the contact rate (λ^w and λ^e) and the matching probability (μ^w and μ^e). Effects are not sensitive to specifications: the Uncorrected Models, the partially-corrected models and the Corrected Models produce very similar results, and hence we only report the results from the Corrected Models.

According to the partial equilibrium search model, a job-seeker characteristic x which increases the interview rate ($\partial\lambda^w/\partial x > 0$) is expected to decrease the matching probability ($\partial\mu^w/\partial x < 0$). The intuition is that, holding the offer distribution constant, an increase in the interview rate increases the expected return to search, and hence increases the reservation wage of the job-seeker and reduces the acceptance probability (see Petrongolo and Pissarides, 2006, for a more detailed discussion). It is an empirical question which effect dominates for a specific characteristic. Table 3.4 shows that the impact on interview rate often takes the opposite direction as the impact on the matching probability, although this is not always the case. Males have a higher matching hazard of 0.069 (0.037). This arises because males have a lower contact rate −0.096 (0.019)

but a higher matching probability 0.059 (0.035).[23] Job-seekers with the highest grade have a significantly higher interview hazard compared to those with lowest grade 0.637 (0.036), yet a significantly lower matching probability −0.163 (0.066). The contact effect dominates, as the combined effect on the matching hazard is positive 0.341 (0.067). For other grade levels the contact rates are also significantly higher than for the lowest grade, the difference in matching probability is insignificant, and the combined effect is a higher matching hazard.

Ethnic minorities have a much lower hazard rate than whites −0.809 (0.108). Indeed, the effect of ethnic minority is more than twice as large as the difference between the highest and lowest educational groups. Interestingly, this effect comes mainly through the matching probability −0.634 (0.106). This is likely to be an employer effect: a lower contact hazard implies that ethnic minorities have a lower reservation wage, yet their matching probability is lower, which implies employers are less likely to recruit ethnic minorities conditional on an interview, which could arise from discrimination, as discussed in Section 3.2. A similar effect is observed for qualified for additional government funding (most likely due to educational or health difficulties).

The results from employer characteristics are less clear-cut. For example, written applications receive more applications 0.135(0.038), but there is a large negative effect on the matching probability −1.085(0.065). This suggests that vacancies which require a written application have a much higher reservation productivity, despite that a written application should convey more information. In contrast, non-manual and skilled vacancies have a lower interview rate by −0.273 (0.030) and −0.130 (0.031), and effects on matching probability is insignificant for skilled vacancies, and is large

[23]The decomposition does not work for some of the variables, due to the fact that we have multiple covariates and that the random effects for the contact function and matching probability are estimated independently. The current discussion is qualitative. Strictly speaking we should exponentiate the coefficients to generate an impact on monthly hazards; coefficients are approximately the estimated percentage difference between the base and the comparison group when the coefficients are small.

Table 3.4 Estimated effect of covariates on hazard rate, contact rate and matching probability

(a) Job-seeker panel	h^w		λ^w		μ^w	
Male	0.069	(0.037)	−0.096	(0.019)	0.059	(0.035)
Exam performance 2	0.148	(0.057)	0.242	(0.031)	−0.041	(0.055)
Exam performance 3	0.321	(0.064)	0.382	(0.035)	−0.031	(0.061)
Exam performance 4 (High)[a]	0.341	(0.067)	0.637	(0.036)	−0.163	(0.066)
Exam performance missing	0.083	(0.059)	0.498	(0.031)	−0.257	(0.056)
Ethnic minority[b]	−0.809	(0.108)	−0.288	(0.043)	−0.634	(0.106)
Additional funding[c]	−0.686	(0.073)	−0.278	(0.033)	−0.506	(0.070)
Obs.	166,625		166,625		36,451	
(b) Vacancy panel	h^e		λ^e		μ^e	
Skilled vacancy[d]	−0.203	(0.062)	−0.273	(0.030)	0.074	(0.140)
Non-manual vacancy[e]	−0.454	(0.062)	−0.130	(0.031)	−0.255	(0.092)
Written application[f]	−1.074	(0.079)	0.135	(0.038)	−1.085	(0.065)
6–10 employees	0.038	(0.077)	−0.019	(0.040)	0.057	(0.062)
11–30 employees	0.115	(0.085)	0.006	(0.043)	0.139	(0.092)
31–50 employees	−0.106	(0.109)	−0.072	(0.058)	0.074	(0.104)
51–100 employees	0.248	(0.110)	0.018	(0.061)	0.210	(0.139)
101–500 employees	−0.022	(0.107)	−0.102	(0.054)	0.158	(0.090)
500+ employees	−0.052	(0.148)	−0.049	(0.082)	0.029	(0.213)
Firm size missing[g]	−0.177	(0.088)	−0.221	(0.049)	0.277	(0.153)
Obs.	26,764		26,764		36,451	

[a] High exam performance consists of school-leavers who achieved four or more high grades in the UK school-leaving exams taken at the age of 16 (O-levels). These school-leavers would generally be qualified to continue onto to take university entrance exams at 18 (A-levels). The base group comprises school-leavers with no exam qualifications or very low grades.

[b] Based on the perceived cultural or ethnic background of the school-leaver.

[c] Refers to school-leavers who qualify for additional training funding as a result of some educational or health difficulties.

[d] Skilled vacancies based on Matching Service occupational classifications.

[e] Non-manual vacancies based on Matching Service occupational classifications.

[f] Vacancies which require a written application.

[g] The base group for firm size comprises employers with ≤ 5 employees.

and negative for non-manual vacancies -0.255 (0.092). The impact of firm size on the contact hazard is generally insignificant, apart from the result that employers do not provide size information in their profile have a lower contact rate -0.221 (0.049). The matching probability does not change systematically with firm size, and the combined effect on the matching hazard also tends to be insignificant. The interpretation is that firm size does not have a one-way relationship with the quality of the vacancy: large firms could have attractive vacancies with training and promotion opportunities and ones which are less attractive than a vacancy at a small local firm.

3.6 Conclusions

In this paper we provide the first estimates of the matching process, using micro data from both sides of the same, well-defined labour market. Our data documents the search spells of both job-seekers and vacancies, and the meetings and matches during the spells. This allows us to investigate the mechanisms by which market size (U and V) and labour market tightness affect the matching process, the causes of the decline of the hazard rate, and the effects of job-seeker and vacancy characteristics.

The fact that we observe agent-level data for both sides of the market allows us to identify the arguments of the matching function from two approaches, one using job-seeker data and the other using vacancy data. This means that biases caused by omitted variables (if we ignore the temporal aggregation) or measurement errors (if we misspecify the spatial aggregation) will manifest themselves in the form of a gap between the two estimates. Estimates of friction and congestion, and the resulting returns to scale, are much more sensitive to the temporal aggregation and cross-market spill-overs, compared to the the results on duration dependence and characteristics. In particular, the effects of temporal aggregation is sizable: the return to scale is much larger before

we correct for the temporal aggregation bias. Using a specification which corrects for both problems, our estimates of the matching function are much closer together than a baseline specification without correcting for those problems. We note that these problems have been overlooked in much of the prior empirical literature.

The hazard rate estimates are quite robust to estimation method. In fact, controlling for observable characteristics and unobserved heterogeneity have very little effect on the pattern of hazards. This finding is quite distinct from most of the literature: Machin and Manning (1999) review a large number of early empirical studies in European countries, and conclude that in most countries (apart from the UK), little evidence is reported for a true negative duration dependence; most of those studies do not control for unobserved characteristics yet the negative duration dependence is not even robust to controls for a few observables. One potential explanation is that the job-seekers and vacancies in our sample are a lot more homogeneous than the ones investigated in other studies, which makes the control of agent characteristics less important. In our sample, on both sides of the market, the decline in matching hazards appear to be completely driven by the fall in contact hazards over the search duration.

These findings have important policy implications. A lot of labour market policies are centered on the issue of reducing long-term unemployment. Our findings contribute to the debate over whether the state should try to force the reservation wages down. The duration pattern of matching probabilities (and hence the reservation wages) is close to constant whereas the contact hazards fall immediately after the first month – the reservation wages will have to fall a lot (yet still might not be adequate) to fully cancel the effects of contact hazards. A related question is whether increasing the contacts between job-seekers and employers (e.g. through forcing the benefit claimants to have certain number of contacts with potential employers) can reduce unemployment. This will depend on the cause of the fall in the contact hazards, or more specifically, whether

the number of contacts fall due to lack of potentially suitable agents in the market, which we explore in the next chapter.

In terms of the effects of job-seeker characteristics, we find that males are less likely to get interviews but are more likely to match conditional on an interview. Higher qualifications leads to more interviews but an interview is less likely to lead to a match. The ethnic minorities, and those on additional subsidies due to learning difficulties/disabilities do worse on both getting interviews and matching conditional on an interview. The results on the matching probabilities highlight the issue of discrimination in the youth labour market. The disadvantaged groups (ethnic minorities and disabled) are likely to have a lower reservation wage due to the lower contact rates which should predict a higher matching probability. Yet their matching probabilities are still lower, which can only attributed to the employer's preferences against hiring from those groups.

Appendix 3.A Temporal aggregation bias: Simulated estimations

This simulation investigates the size of temporal aggregation bias, and tests whether our estimation corrects for this bias. We generate 10-year daily data using the following process. The job-seeker and vacancy inflows $\{u, v\}$ on each day τ are generated by Poisson process

$$u_\tau \sim Poisson(u_n), \qquad v_\tau \sim Poisson(v_n),$$

where $u_n = 5.00$ in year 1, 3, 5, 7 and 9, and 5.10 in year 2, 4, 6, 8 and 10, $v_n = 5.05$ in year 1, 3, 5, 7 and 9, and 5.00 in year 2, 4, 6, 8 and 10. The fluctuating inflow rates generate cyclical trends in the U, V stocks. The average inflow of U is slightly higher than the average inflow of V, which creates a slack market where the average U stock is higher than the V stock.

The daily contacts and matches are specified as

$$C_\tau \sim \text{Poisson}(A_\tau U_\tau^\alpha V_\tau^\beta), \qquad M_\tau \sim \text{Binomial}(C_\tau, \mu),$$

where $A_\tau = e^{-1+\varepsilon_\tau}$, $\alpha = \beta = 0.5$, $\mu = 0.1$, $\varepsilon_\tau \sim \text{Normal}(0, \sigma)$, $\sigma = 0$, 0.1, 0.35, 0.5, and 0.7 in the five sets of simulations, each with 300 replications.

Matches are the only source of outflows, and stocks changes daily as a result of inflows and outflows:

$$U_\tau = U_{\tau-1} - M_{\tau-1} + u_{\tau-1}, \qquad V_\tau = V_{\tau-1} - M_{\tau-1} + v_{\tau-1}.$$

The data is generated for 3,600 days, and is aggregated into a 120-month (30 days per month) time series data, with the monthly total contacts and matches $\{C_t, M_t\}$, the

Table 3.A1 Simulated contact function estimation (300 replications)

Truth: $\alpha = \beta = 0.5$

	Standard deviation of ε				
	0.001	0.1	0.35	0.5	0.7
(a) Daily Estimates					
$\hat{\alpha}$	0.502	0.503	0.508	0.509	0.509
	(0.017)	(0.020)	(0.039)	(0.048)	(0.055)
$\hat{\beta}$	0.503	0.502	0.507	0.508	0.511
	(0.016)	(0.020)	(0.038)	(0.049)	(0.049)
(b) Uncorrected Model					
$\hat{\alpha}$	0.286	0.283	0.278	0.272	0.247
	(0.042)	(0.041)	(0.050)	(0.046)	(0.052)
$\hat{\beta}$	0.283	0.280	0.274	0.269	0.245
	(0.041)	(0.039)	(0.043)	(0.047)	(0.054)
(c) Corrected Model					
$\hat{\alpha}$	0.474	0.466	0.466	0.479	0.471
	(0.101)	(0.105)	(0.120)	(0.125)	(0.168)
$\hat{\beta}$	0.518	0.515	0.520	0.515	0.514
	(0.095)	(0.099)	(0.118)	(0.133)	(0.160)

beginning-of-month stocks $\{U_t, V_t\}$, and the total inflows over the month $\{u_t, v_t\}$. The first 12 months are excluded from the estimates because the stocks are building up towards the equilibrium level.

Panel (a) reports estimates of aggregate contacts using the daily data with the Poisson link, which is expected to recover the true coefficients. We then conduct estimates using two different specifications (both also include a monthly trend): one with only the beginning-of-month stocks (as in our Uncorrected Model, reported in Panel (b)), one with both the stocks and inflows (as in our Corrected Model, reported in Panel (c)). For the Corrected Model $\hat{\alpha}$ is the sum of the coefficients of $\log U_t$ and $\log u_t$ and $\hat{\beta}$ is the sum of the coefficients of $\log V_t$ and $\log v_t$.

The daily estimates correctly recovers the true parameters. The Uncorrected Model estimates are severely biased towards zero, with the size of bias mildly increases with σ, the variance of error in the contact technology. The Corrected estimates appear to be unbiased to the second digit for all σ's, which suggests that our method corrects for the temporal aggregation bias.

Appendix 3.B Map of Lancashire

Fig. 3.B1 Lancashire consists of 14 labour markets, which are categorized into three neighbourhood districts.



Appendix 3.C Additional descriptive statistics

Table 3.C1 Summary of raw data: hazards, contacts and matching probabilities

(a) Job-seekers						
n	$\sum_i M_{in}$	$\sum_i C_{in}$	$\sum_i \tau_{in}$	$\hat{h}^w_n \times 30$	$\hat{\lambda}^w_n \times 30$	$\hat{\mu}^w_n \times 100$
1	579	16,614	899,372	0.020	0.562	3.485
2	1,018	7,891	987,429	0.031	0.243	12.901
3	584	4,546	691,290	0.026	0.200	12.846
4	399	2,927	494,721	0.025	0.180	13.632
5	215	1,592	286,953	0.023	0.169	13.505
6	127	950	204,684	0.019	0.141	13.368
7	59	562	153,252	0.012	0.111	10.498
8	44	405	114,977	0.012	0.107	10.864
9	22	256	81,887	0.008	0.095	8.594
10	13	207	60,349	0.007	0.104	6.280
11/12	19	256	85,596	0.007	0.091	7.422
13+	16	245	114,348	0.004	0.065	6.531
Total	3,095	36,451	4,174,858	0.023	0.265	8.491
(b) Vacancies						
n	$\sum_j M_{jn}$	$\sum_j C_{jn}$	$\sum_j \tau_{jn}$	$\hat{h}^e_n \times 30$	$\hat{\lambda}^e_n \times 30$	$\hat{\mu}^e_n \times 100$
1	2,062	23,238	226,055	0.277	3.125	8.873
2	593	7,107	243,906	0.074	0.886	8.344
3	115	1,772	152,945	0.023	0.352	6.490
4	68	1,979	127,367	0.016	0.472	3.436
5	78	641	95,532	0.025	0.204	12.168
6	33	434	73,098	0.014	0.180	7.604
7	22	349	67,307	0.010	0.158	6.304
8	36	257	54,403	0.020	0.144	14.008
9	29	160	42,404	0.021	0.115	18.125
10	13	109	30,108	0.013	0.110	11.927
11/12	20	140	50,536	0.012	0.084	14.286
13+	26	265	116,059	0.007	0.069	9.811
Total	3,095	36,451	1,279,720	0.074	0.866	8.491

[1] The number of observations with $n = 1$ is slightly smaller than the number of spells, and the difference is due to left-censored 22 job-seeker spells and 615 vacancy spells. Prior to the sampling period, the exposure is observed but any contacts are unobserved. The left-censored spells and their contacts and matches, are included in the sample summarized above, and included in the analysis. For instance, if an agent has already entered the market for two months at the start of the sampling period, their first record enters the summary of $n = 3$.

Table 3.C2 Observable characteristics by results of job-seeker search spell. Job-seekers and vacancies who match within the system have similar observable characteristics to those who match outside the system.

(a) Job-seekers

	N	Male	Ethnic minority	Receiving subsidy	Higher-level qualifications[a]
Censored	35,882	0.58	0.07	0.08	0.31
Match within system	3,095	0.60	0.03	0.03	0.29
Match outside system	10,116	0.57	0.04	0.04	0.27
Total	49,093	0.58	0.06	0.07	0.30

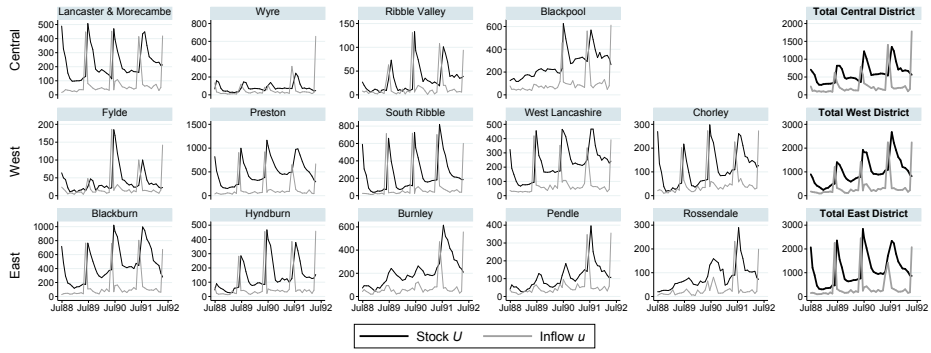
(b) Vacancies

	N	Skilled	Non-manual	Written application	Median or large firms[b]
Censored	7,093	0.59	0.61	0.38	0.19
Match within system	3,095	0.48	0.46	0.19	0.22
Match outside system	7,322	0.51	0.49	0.21	0.26
Total	17,510	0.54	0.54	0.27	0.22

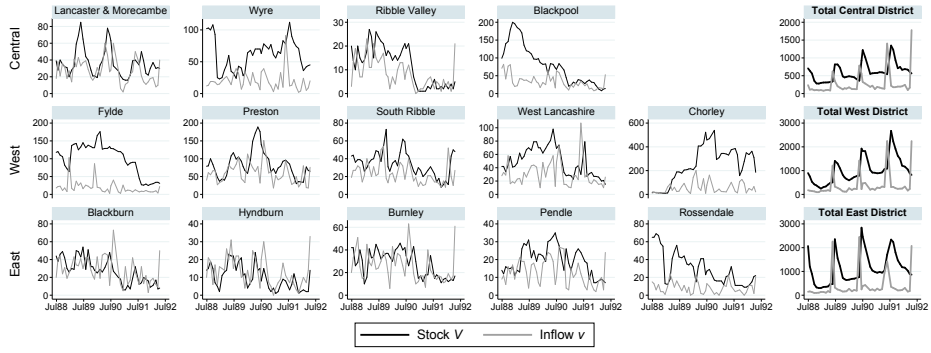
[a] With more than three O-level or A-level qualifications.
[b] Firms with more than 50 employees.

Fig. 3.C1 Monthly time-series of stock, inflow and labour-market tightness by local labour market, $\{U_{kt}, V_{kt}, u_{kt}, v_{kt}\}$ and V_{kt}/U_{kt}. Each row displays the markets within the same region, and the last column displays the variables by region, $\{U_{rt}, V_{rt}, u_{rt}, v_{rt}\}$, and V_{rt}/U_{rt}.

(a) Stock U and inflow u



(b) Stock V and inflow v



(c) Labour-market tightness V/U



The horizontal line indicates the median tightness for all 630 district-months.

Appendix 3.D Additional regression results

Table 3.D1 Estimates of Corrected Models: raw coefficients and spill-over parameters

	Hazards		Contacts		Matching Probabilities	
	h^w	h^e	λ^w	λ^e	μ^w	μ^e
π_{11} $\log U_{rt}$	-0.214	0.428	-0.108	0.411	-0.089	0.054
	(0.074)	(0.090)	(0.032)	(0.047)	(0.070)	(0.125)
π_{12} U_{kt}/U_{rt}	-0.956	0.964	-0.127	1.027	-0.628	-0.143
	(0.429)	(0.503)	(0.180)	(0.312)	(0.406)	(0.525)
π_{13} $\log V_{rt}$	0.388	0.021	0.289	-0.207	0.068	0.079
	(0.113)	(0.136)	(0.046)	(0.086)	(0.108)	(0.121)
π_{14} V_{kt}/V_{rt}	0.726	-0.182	-0.096	0.406	0.405	0.302
	(0.448)	(0.519)	(0.185)	(0.326)	(0.420)	(0.468)
π_{21} $\log u_{rt}$	-0.315	0.060	-0.219	0.445	-0.112	-0.274
	(0.045)	(0.050)	(0.018)	(0.026)	(0.042)	(0.046)
π_{22} u_{kt}/u_{rt}	-0.749	-0.173	0.174	0.564	-0.375	-0.831
	(0.322)	(0.324)	(0.127)	(0.197)	(0.274)	(0.338)
π_{23} $\log v_{rt}$	0.166	-0.321	0.095	-0.402	0.019	0.015
	(0.056)	(0.070)	(0.023)	(0.043)	(0.054)	(0.062)
π_{24} v_{kt}/v_{rt}	1.165	-0.132	0.879	-0.399	0.270	0.344
	(0.247)	(0.279)	(0.100)	(0.166)	(0.241)	(0.231)
δ^w	0.237	0.382	1.167	0.350	0.167	0.184
	(0.065)	(0.197)	(1.085)	(0.060)	(0.086)	(0.111)
δ^e	0.227	0.488	0.329	1.011	0.114	0.128
	(0.062)	(0.515)	(0.068)	(0.672)	(0.154)	(0.177)
α	0.471	0.488	0.673	0.856	-0.201	-0.220
	(0.095)	(0.112)	(0.040)	(0.059)	(0.090)	(0.141)
β	0.554	0.700	0.384	0.391	0.087	0.094
	(0.122)	(0.147)	(0.051)	(0.097)	(0.118)	(0.130)
$\alpha+\beta$	1.025	1.189	1.057	1.247	-0.114	-0.126
	(0.139)	(0.167)	(0.059)	(0.099)	(0.137)	(0.184)
Log Likelihood	-16041	-9382	-89291	-37246	-9839	-9807
Obs.	166625	26764	166625	26764	36451	36451

Notes: This table reports the raw coefficients for the "Corrected" columns in Table 3.3 (estimates of Equation (3.27) and its companions).

Chapter 4

Evidence on stock-flow search

4.1 Introduction

The aggregate matching function, typically taking the Cobb-Douglas form, is a widely-used tool for modelling labour market frictions. It summarizes the number of matches as a function of total job-seekers and vacancies. It is assumed to be concave and increasing in both arguments, such that an increase in the number of searching agents on either side of the market yield more matches with decreasing returns. It is also typically assumed to have constant returns to scale, following the early empirical work of Pissarides et al. (1986) and theoretical discussion of Diamond (1982). However, despite its empirical usefulness, the Cobb-Douglas matching function does not have a micro-foundation, and the mechanism that drives the aggregate phenomenon described by the matching function remains as something of a "black box" (Petrongolo and Pissarides, 2001).

Although the frictions which result in the aggregate matching function are not precisely defined, they are often assumed to result from the fact that agents on one side of the market (e.g. job-seekers) are not fully informed about suitable partners on the other side (e.g. vacancies). Suitable pairs of job-seeker and vacancy exist in the market, yet they do not form matches immediately because agents do not immediately observe all the trading opportunities available (Coles and Smith, 1998; Gregg and Petrongolo, 2005). The market is assumed to be sufficiently large relative to the speed with which agents meet each other such that agents never run out of potential partners, and hence the rate at which agents meet and match is invariant to how much time they have spent in the market. This is plausible if the process by which agents acquire information about and/or meet with potential partners is costly in terms of time and effort.

However, in an environment where there is a database of potential partners, such frictions may become much less important. Instead, frictions may arise because of co-ordination failures, or simply because of a lack of suitable partners. In such an environment, a natural explanation for frictions is provided by the stock-flow matching

model (Coles and Muthoo, 1998; Coles and Smith, 1998). In the stock-flow model the meeting friction is assumed away. Agents can search the other side of market immediately through some centralized agency or marketplace.[1] This frictionless meeting process implies that matches will only occur between new entrants (the "flow") on one side of the market and those on the other side who have not found a partner when they were new (the "stock").

To explain the coexistence of unemployed job-seekers and vacancies, the stock-flow model assumes some degree of heterogeneity in terms of geographic location, skill level or occupation, which cannot be costlessly changed. As a result, a suitable trading partner may not readily exist in the market, and those who do not find a match immediately have to wait for new opportunities to arise. This implies that exit rates from search will exhibit strong duration dependence. New job-seekers and vacancies have a large pool of potential partners, and therefore have a high exit rate; job-seekers and vacancies which fail to match with the existing stock have a much smaller pool of potential partners and therefore have a low exit rate.

Previous studies have provided evidence that stock-flow matching is more successful in describing the pattern of matches observed at both the aggregate and the individual level (Alvarez de Toledo et al., 2011; Andrews et al., 2013; Coles et al., 2010; Coles and Petrongolo, 2008; Gregg and Petrongolo, 2005; Kettemann et al., 2017; Kuo and Smith, 2009; Sasaki, 2008; Soininen, 2007). However, the key distinction of the stock-flow matching model originates from its assumptions about the process by which agents *contact* each other. Therefore, a more direct test of stock-flow matching requires an analysis of the contacts between job-seekers and vacancies. A systematic study of contacts for both job-seekers and vacancies is also useful in comparing the stock-flow

[1] The presence of centralized agency, which allows employers to advertise vacancies (in addition private private advertisements), by itself does not rule out random search, according to the discussion of Pissarides (1979).

model with other mechanisms which also predict negative duration dependence of exit rates (e.g. the ranking model of Blanchard and Diamond, 1994).

Using unique data which contain information on contacts (i.e. interviews) in a particular labour market collected by the Matching Service for Lancashire, we test the assumptions and implications of the stock-flow matching model. Our descriptive evidence shows that the frequency of contacts decreases with search duration, and that "old" agents (those who have been in the market for some time) are far more likely to contact "new" agents than old ones on the other side of the market. Such findings are consistent with the stock-flow model. Our econometric estimates show that it is the inflow rate of new agents, rather than the total stock of agents, which determines the contact rates of existing agents, controlling for unobserved heterogeneity.

Our paper follows and builds on Andrews et al. (2013), who test the stock-flow matching theory, but only consider the pattern of matches. In this paper, we draw on the same data source,[2] and we extend their framework to examine whether contacts follow a stock-flow pattern and thus provide unique insights into the "black box" of the matching function.

The paper is organized as the follows: Section 4.2 reviews the empirical literature testing the stock-flow matching model against random matching model, and the limited number of other studies which examine job-search contacts (as opposed to matches). Section 4.3 discusses the empirical framework. Section 4.4 introduces the features of the data. Some descriptive evidence is presented in Section 4.5, followed by the estimation results in Section 4.6. Section 4.7 concludes.

[2] The data organisations are slightly different, although the structures are similar. Part of the work of this thesis is to write Stata codes which generate the estimation panels from the raw data.

4.2 Previous literature

4.2.1 Tests of stock-flow matching model

Coles and Smith (1998) show that the stock-flow model differs from random matching models in terms of three testable predictions. First, the random matching model predicts that matches between stocks are as likely as matches between stock and flow. In the stock-flow model, in contrast, matches between stocks never occur. Second, the stock-flow model predicts that the hazard rate is initially high and then becomes low. Third, under random matching the hazard rate depends on the *same* variables (i.e. the aggregate stocks of job-seekers and vacancies) over the duration of search, while with stock-flow matching, the hazard rate depends on different variables at different durations: initially on the total stock on the other side of the market, but then only on the new inflows.

Using data collected from UK job centres, Coles and Smith (1998) provide descriptive evidence for stock-flow matching: the stock of vacancies has a strong positive impact on the exit rate of job-seekers in the first two weeks of unemployment, but has no effect on the exit rate of job-seekers who have been unemployed for more than four weeks. Ebrahimy and Shimer (2010) and Shimer (2007) also find that co-movements of aggregate time series of unemployment and vacancy stocks are captured better by the stock-flow model compared to the random model.

More formal tests of stock-flow matching against random matching are provided by a number of studies. Some of them use market-level time series data (Coles and Petrongolo, 2008; Gregg and Petrongolo, 2005; Sasaki, 2008; Soininen, 2007) and others use data of individual unemployment (and/or vacancy) spells (Alvarez de Toledo et al., 2011; Andrews et al., 2013; Kettemann et al., 2017; Kuo and Smith, 2009). Most of them adopt a flexible structure which encompasses both random matching and stock-flow matching mechanism. In general they conclude that the stock-flow matching

model outperforms the random matching model in capturing the observed matches or unemployment outflows.

4.2.2 Evidence from search activities

This paper provides another direct empirical test of random versus stock-flow matching, not by looking into the outcomes of search, but the search process itself.

Due to lack of data availability, there is only a small number of early empirical studies which track search activities over the unemployment spell to determine the true duration dependence, i.e. the causal impact of duration on search, rather than any selection effects. For example, Erens and Hedges (1990) find evidence of a moderate decline of the number of hours spent looking for work over a spell of unemployment, and Schmitt and Wadsworth (1993) suggest that the long-term unemployed search for work less intensively.

More recently, using high-frequency regional survey data on search activities of unemployed job-seekers in the US recession of 2009–10, Krueger and Mueller (2011) find that unemployment duration has a significant negative impact on the amount of time spent on search each week, after controlling for person fixed effects and changes in unemployment benefits. Also, there is little evidence that search time increases upon the exhaustion or expiration of unemployment insurance. The stock-flow matching model offers a consistent explanation to such findings (i.e. the key reason that some agents remain unmatched is a lack of suitable partners, rather than informational friction). However, as noted in the paper, there are also other explanations. For example, job-seekers may search more efficiently as they accumulate experience, and human capital depreciation may reduce the expected wage and hence the amount of search activities.

Developments in search technology, especially the use of the internet as an dominant method of job-search, significantly lowers the time and monetary cost of search.[3] Faberman and Kudlyak (2014) study search and application behaviour using a large dataset of applications collected from an online job search engine which matches hourly-paid jobs with job-seekers who tend to be lower-skilled and less-educated. Their analysis provides two major implications. First, controlling for observed and unobserved heterogeneity in characteristics, the number of applications significantly decreases over the duration of search.[4] Second, people with longer search spells tend to send out a larger number of applications per unit of time. However, the findings are only suggestive, because labour market status and actual search duration are unobserved from their data, only approximated by the observed applying activities. It may be that individuals who search longer are more intensive users of the particular search engine relative to other users.

This paper is distinguished from the previous studies on search activities because we focus on testing the predictions of stock-flow against those of random matching mechanism. In addition, our data also identify search durations of job-seekers and vacancies, and the individuals who contact each other.

4.3 Empirical Framework

We generalise the discrete-time version of the stock-flow matching function introduced by Andrews et al. (2013) into a stock-flow contact function. According to the discussion of Pissarides (2000, Chapter 6), a contact is a meeting between a job-seeker and an employer with vacancy (e.g. an application/interview). Only a fraction of contacts become matches. This is because, although job-seekers and vacancies are *ex ante*

[3] However, a pioneering study by Kuhn and Skuterud (2004) finds no evidence that job-seekers who use the internet have shorter unemployment durations. A follow-up study, Kuhn and Mansour (2014), for a later period finds that internet job search to be more effective.

[4] This finding is closely related to the stock-flow matching model, and is replicated on our data (see more details in Section 4.5).

homogeneous, the productivity of each pair is heterogeneous and unknown until they meet, and a pair will only match if their revealed productivity exceeds the reservation productivity.

Assume that within a "window" period of finite length, new agents can screen the entire pool of potential partners which exist upon their entry. Denote the total stock of job-seekers and vacancies in the market by U and V. The window period can be used to split the total stock into two components: the "new" stock, denoted by u and v, which include those who are still in their window period screening the existing opportunities, and the "old" stock, denoted by $\bar{U} \equiv U - u$ and $\bar{V} \equiv V - v$, which are those who have finished their screening period. It follows that contacts are split into four categories, depending on whether the job-seeker and vacancy is new or old, denoted by C_{11}, C_{12}, C_{21} and C_{22}, where the subscripts indicate the elapsed search durations of the pair (1 for "new" and 2 for "old"), the first place for job-seeker and the second place for vacancy.

If contacts occur randomly, the number of contacts between new and old agents is simply the relevant proportion of the total number of contacts, C. If the total number of contacts follows a Cobb-Douglas function,

$$C = C(U,V) = U^{\alpha}V^{\beta},$$

then, for example, the number of contacts between new job-seekers and old vacancies, C_{12}, can be written as

$$C_{12} = \frac{u}{U}\frac{\bar{V}}{V}C = u\bar{V}U^{\alpha-1}V^{\beta-1}, \tag{4.1}$$

where $\frac{u}{U}$ is the proportion of new job-seekers among all job-seekers, and $\frac{\bar{V}}{V}$ is the proportion of old vacancies among all vacancies, and multiplying the two proportions to the total contacts gives the number of contacts between new job-seekers and old vacancies.

To capture the spirit of stock-flow matching, we introduce additional parameters a_{11}, a_{12}, a_{21} and a_{22} to allow for uneven split of total contacts across the four types:

$$\begin{aligned} C_{11} &= a_{11}\frac{u}{U}\frac{v}{V}C = a_{11}uvU^{\alpha-1}V^{\beta-1}, \\ C_{12} &= a_{12}\frac{u}{U}\frac{\bar{V}}{V}C = a_{12}u\bar{V}U^{\alpha-1}V^{\beta-1}, \\ C_{21} &= a_{21}\frac{\bar{U}}{U}\frac{v}{V}C = a_{21}\bar{U}vU^{\alpha-1}V^{\beta-1}, \\ C_{22} &= a_{22}\frac{\bar{U}}{U}\frac{\bar{V}}{V}C = a_{22}\bar{U}\bar{V}U^{\alpha-1}V^{\beta-1}. \end{aligned} \tag{4.2}$$

Under random matching $a_{11} = a_{12} = a_{21} = a_{22} = a$. Under "pure" stock-flow matching, $a_{22} = 0$ because, given the meeting technology, any potentially suitable partners would already have been interviewed in the initial period. More generally, we suppose that contacts between old job-seekers and old vacancies are possible[5] but unlikely, which gives $a_{12} > a_{22}$ and $a_{21} > a_{22}$. We focus on testing these two relationships; the rest of relationships between a's are less clear from the theory.

The contact rate for an individual job-seeker, c^U, follows directly from (4.2), and varies with the search duration of the job-seeker. A new job-seeker has a logged contact rate written as

$$\log c^U_{1.} = \log\left[c^U_{11} + c^U_{12}\right] = \log[a_{11}v + a_{12}\bar{V}] + (\alpha-1)\log U + (\beta-1)\log V, \tag{4.3}$$

while for old job-seekers, the logged contact rate is

$$\log c^U_{2.} = \log\left[c^U_{21} + c^U_{22}\right] = \log[a_{21}v + a_{22}\bar{V}] + (\alpha-1)\log U + (\beta-1)\log V. \tag{4.4}$$

[5]For example, if agents have heterogeneous window periods, or if some agents change their preferences during search.

Under random matching these contact rates are equal because $a_{11} = a_{12} = a_{21} = a_{22}$. This implies that being old or new (and how long an agent has been searching in general) is irrelevant for the contact rates. Under stock-flow matching, since a_{22} is small, the contact rate of job-seekers is predicted to fall sharply if they fail to match within the initial period. Similar expressions define the contact rate for new and old vacancies, $c^V_{.1}$ and $c^V_{.2}$.

It is convenient to estimate a linear version of (4.3) and (4.4). For new job-seekers, this is

$$\log c^U_{1.} = \rho_0 + \rho_1 \log V + \rho_2 \log v + \rho_3 \log U + \rho_4 \log u + \varepsilon^U_{1.}, \tag{4.5}$$

and for old job-seekers this is

$$\log c^U_{2.} = \pi_0 + \pi_1 \log V + \pi_2 \log v + \pi_3 \log U + \pi_4 \log u + \varepsilon^U_{2.}. \tag{4.6}$$

Equation (4.4) implies that

$$\frac{\partial \log c^U_{2.}}{\partial \log V} = \frac{a_{22}V}{a_{21}v + a_{22}\bar{V}} + \beta - 1 \equiv \pi_1, \qquad \frac{\partial \log c^U_{2.}}{\partial \log v} = \frac{(a_{21} - a_{22})v}{a_{21}v + a_{22}\bar{V}} \equiv \pi_2,$$
$$\frac{\partial \log c^U_{2.}}{\partial \log U} = \alpha - 1 \equiv \pi_3, \qquad \frac{\partial \log c^U_{2.}}{\partial \log u} = 0 \equiv \pi_4.$$

Therefore π_1, π_2 and π_3 uniquely identify α, β and a_{22}/a_{21}. The estimate of α is simply $\pi_3 + 1$.[6] To obtain the estimate of β we add together the effects of the old and new vacancy stocks, i.e. $\beta = \pi_1 + \pi_2$. And a_{22}/a_{21} derived from $\partial \log c^U_{2.}/\partial \log v$ can be written as

$$\frac{a_{22}}{a_{21}} = \frac{v}{V(1-\pi_2)^{-1} - \bar{V}}. \tag{4.7}$$

[6]Empirically it is possible to have $\hat{\pi}_4 \neq 0$. Therefore we estimate $\hat{\alpha} = \hat{\pi}_3 + \hat{\pi}_4 + 1$ for the full impact of the change in the total unemployment stocks.

Therefore a positive π_2 implies that $a_{22}/a_{21} < 1$. From (4.6), a positive estimate of π_2 means that the flow of new vacancies has a significant effect on the contact rate of old job-seekers over and above the effect coming from the stock of all vacancies. Analogous arguments apply to the contact rate for old vacancies, c^V:

$$\log c^V_{.2} = \pi_0 + \pi_1 \log V + \pi_2 \log v + \pi_3 \log U + \pi_4 \log u + \varepsilon^V_{2.}. \tag{4.8}$$

In this model, a positive π_4 indicates that the flow of new job-seekers has a significant effect on the contact rate of old vacancies over and above the effect coming from the stock of all job-seekers.

To model monthly contacts, we assume that the number of contacts in each month of search is Poisson distributed. For job seeker i in month t, this is written as

$$c^U_{it} = E_{it} \exp(\lambda^U_{it}),$$

where E_{it} is the exposure in days within month t and λ^U is the Poisson parameter. Under random matching, the mean of the Poisson process is a linear function of the logged beginning-of-month stocks:

$$\lambda^U_{it} = (\alpha - 1) \log U_t + \beta \log V_t + e_i,$$

where e_i is the job-seeker spell level random-effect following a log-gamma distribution, and is assumed to be uncorrelated with $\log U$ and $\log V$ (relaxation of this assumption is briefly discussed in the Section 4.6.3).

Under stock-flow matching, from (4.5) and (4.6) the mean of the Poisson process is a function of new and old stocks, and the effect of those stocks varies with elapsed duration. Let $x_{it} = [\log V_{it}, \log v_{it}, \log U_{it}, \log u_{it}]$, then the Poisson parameter is modeled

by:

$$\lambda_{it}^{U} = d_{it}x'_{it}\rho + (1-d_{it})x'_{it}\pi + e_i, \tag{4.9}$$

where d_{it} is a dummy variable indicating whether job-seeker i in month t is new. The parameter vector ρ therefore contains estimates of the effect of new and old stocks on new job-seekers ($\rho_1,\ldots,\rho_4$ in 4.5), while π contains estimates of the effect of new and old stocks on old job-seekers ($\pi_1,\ldots,\pi_4$ in 4.6). Using this specification we estimate the contact rate, pooling the old and new agents.[7] An exactly symmetric model is used for λ_{jt}^{V}, the contact rate for vacancies. The model is estimated using conditional maximum likelihood method. To control for heteroskedasticity, we bootstrap standard errors for all the estimates.

4.4 Data

The data are from the Matching Service for Lancashire for the period from March 1988 to June 1992. This government-funded Matching Service registers the young, unemployed job-seekers aged 15-18, and provides a platform for employers to advertise their vacancies. Using their computerised records, Matching Service staff facilitate a screening process for the job-seekers and employers, and arrange interview meetings for them. When the job-seekers and employers find a match (either within or outside the system) or decide to stop searching, they notify the Matching Service to close their registration. Therefore our data record interviews and hires, and the dates which approximate the start and end dates of the search spells of the unemployed job-seekers and vacancies.[8]

[7] Unlike in Chapter 3, the models reported in this chapter do not include observable characteristics, which are less important because the key variables are aggregate rather than individual level.

[8] More information on the raw data, and previous studies using the same data source are provided in Chapter 3.

In this paper we focus on the interview contacts. A "contact" has two meanings. First, job-seekers and vacancies can view all agents on the other side of the market, so an initial screening process can occur with very little friction, as in a stock-flow environment. Second, job-seekers send applications to vacancies and interviews take place. This second kind of contact is not frictionless, because interviews take time and there is some limit to the number of interviews which can take place simultaneously. Nevertheless, given the initial screening process we hypothesize that all suitable bilateral meetings can take place within a relatively short period of time.

Since our data record contacts in monthly intervals, we create two panel datasets, one at job-seeker-month and one at vacancy-month level.[9] For each job-seeker/vacancy month the data record the elapsed duration of search, and the number of contacts made within that interval.

The job-seeker panel contains 35,269 unemployed job-seekers who have 44,876 search spells. These unemployed job-seekers make 34,645 contacts to job vacancies registered on the system. The vacancy search panel contains 12,401 vacancies with 17,345 vacancy spells.[10] The average duration of job-seeker and vacancy search spells are 90 and 80 days, and the medians are 62 and 25 days, respectively. The number of contacts per month are summarized in Table 4.1. The average monthly contacts are 0.2 for jobseekers and 0.6 for vacancies.

Although the Matching Service is a free option for all job-seekers and employers, it is not compulsory for them to use it as their primary search channel. In fact, three

[9]The job-seeker and vacancy panels used in this chapter are from an earlier version of data organisation compared to the Chapter 3, which is reason that the sample sizes vary slightly. This earlier version of data uses a slightly different way to identify individual job-seeker spells and the end dates of their search spell, and does not contain agent-level characteristics. Also, in most of this chapter we consider Lancashire as one labour market (apart from the discussion in Section 4.6.3), which is why the stocks are only indexed by time.

[10]There are more vacancy spells than vacancies, which is because 14% of vacancies have multiple positions. A more complete description of multiple positions within vacancies is provided by Andrews et al. (2008).

quarters of the job-seeker spells and slightly over half of the vacancy spells do not have any contacts. We believe that these agents still do some search through the Matching Service and create congestion for the rest of the agents (e.g. by making and receiving applications which do not result in interviews), but it is also plausible that they use the Matching Service less intensively than the rest of agents. Therefore we also investigate the sub-samples which exclude the inactive agents and keep only the ex post sample of agents at risk. On average the active job-seekers and vacancies have 0.5 and 1.8 contacts per month. The average durations are 113 and 46 days for the active job-seekers and vacancies, and the medians are 79 and 21 days.

Table 4.1 Summary statistics of monthly contacts. Active job-seekers and vacancies have at least one job contact over the search spell.

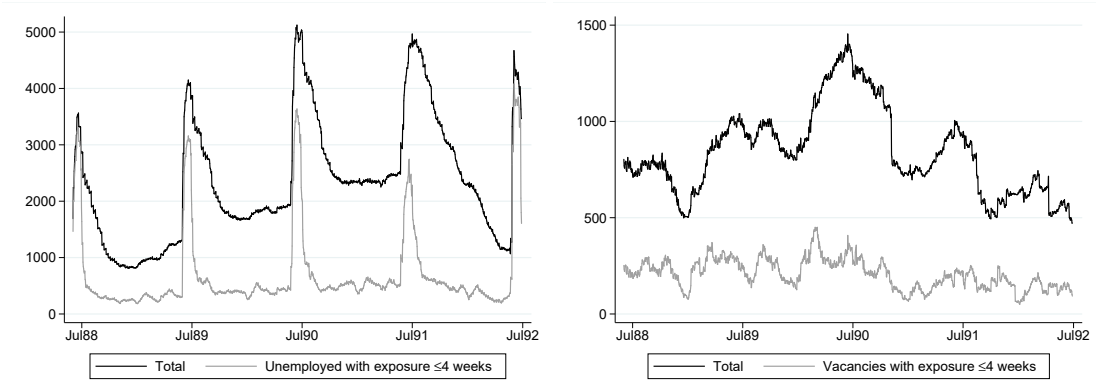
Variable	N	Median	Mean	Std. Dev.	Min.	Max.
Job-seeker panel:						
Monthly contacts - all job-seekers	163,205	0	0.212	0.658	0	15
Monthly contacts - *active* job-seekers	67,169	0	0.516	0.946	0	15
Vacancy panel:						
Monthly contacts - all vacancies	52,880	0	0.611	2.058	0	56
Monthly contacts - *active* vacancies	17,887	1	1.806	3.220	0	56

To estimate (4.9) and its vacancy counterpart we require measures of the stocks of job-seekers and vacancies disaggregated by their elapsed time in the market. To do this we need to decide the length of time it takes for an agent to search the stock on other side of the market. Our baseline model assumes a window period of $k = 4$ weeks, but in robustness checks we test whether our results change for different window periods. The window period defines our new and old agents: on each day τ, the new agents are those in the market who arrived between $\tau - k$ and τ, and the old agents are those in the market who arrived prior to $\tau - k$.

The aggregate series of new and total stocks, u, U, v, V over the sample period are plotted in Figure 4.1. The stocks of job-seekers exhibit a strong seasonal pattern. The

spikes of the job-seeker stocks correspond to the inflow of school-leavers in July each year. The seasonal pattern in vacancy stocks are less strong, although the total stock of vacancies also peaks during summer months. The average stock of new job-seekers is 885, and that of old job-seekers is 1,929. The average stock of new vacancies is 213, and that of old vacancies is 649. The daily stocks are merged to the records in the job-seeker and vacancy spell panels by the beginning dates of each job-seeker and vacancy months.

Fig. 4.1 Total and new stocks of job-seekers and vacancies, $k = 4$ weeks.



(a) Stocks of job-seekers

(b) Stocks of vacancies

To test whether and how contact rates are affected by the presence of inactive agents, we also construct the stocks of active agents, and use them to estimate the contact rates for the sample of active agents. The average active new and old stocks of job-seekers are 325 and 922, and those of vacancies are 168 and 525. The general time-series patterns are similar to those of aggregate stocks reported in Figure 4.1.

4.5 Descriptive Evidence

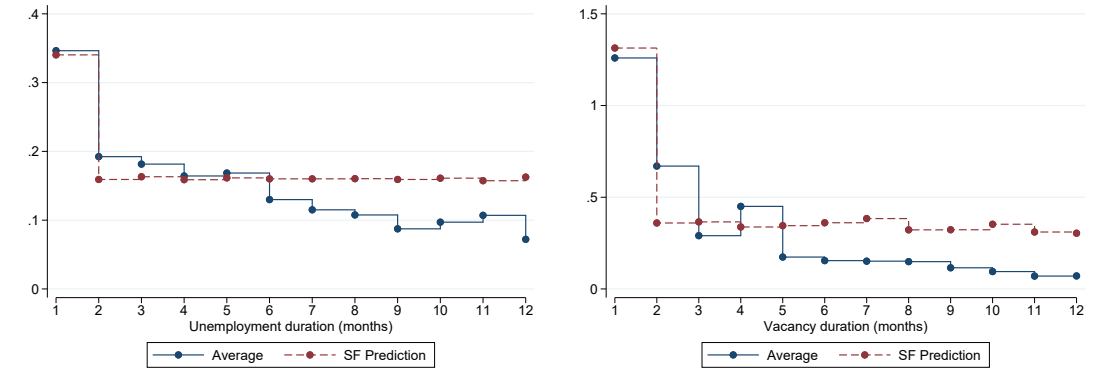
In this section we present evidence that (i) the contact rate of both job-seekers and vacancies exhibits negative duration dependence and (ii) old agents are more likely to contact new rather than old agents on the other side of the market. Both of these facts are

consistent with the predictions of stock-flow matching. However, they are also potentially consistent with random matching with unobserved heterogeneity. Our econometric estimates in Section 4.6 attempt to distinguish between these two hypotheses.

4.5.1 Duration dependence

In Figure 4.2 we summarise the number of contacts for job-seekers and vacancies over the elapsed search months. The solid line in Panel (a) of Figure 4.2 shows that the average number of contacts per month made by job-seekers decreases over the unemployment duration. The solid line in Panel (b) of Figure 4.2 displays the summary for vacancies. Vacancies have more contacts per month compared to job-seekers, especially in the first few months they are posted; and the number decreases as duration increases. The patterns are robust to the exclusion of inactive agents.

Fig. 4.2 Contacts per month of job-seeker and vacancy search.



(a) Job-seekers (b) Vacancies

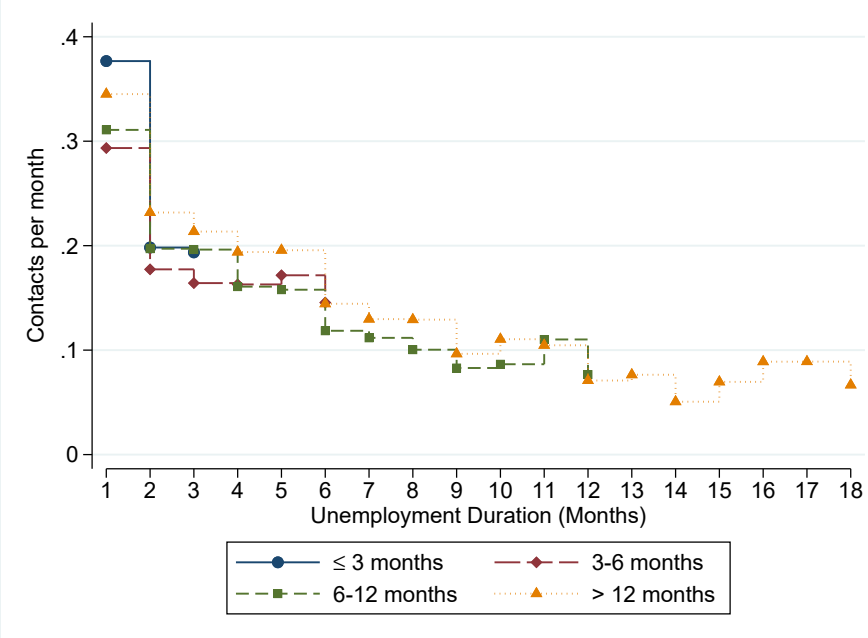
Notes: The stock-flow predictions of the contact rates ("SF Prediction") are estimated using $c_{it}^{U} = \beta_0 \log v_t + \beta_1 \log \bar{V}_t + \beta_2 d_{it} \log \bar{V}_t + \varepsilon_{it}$ and $c_{jt}^{V} = \beta_0 \log u_t + \beta_1 \log \bar{U}_t + \beta_2 d_{jt} \log \bar{U}_t + \varepsilon_{jt}$, for job-seekers and vacancies respectively.

The raw duration dependence observed in Figure 4.2 is consistent with stock-flow matching, and could be explained by the fact that the size of old stocks greatly exceeds the size of new stocks. Hence, once an agent has interviewed all existing suitable partners,

the number of potential partners declines and the contact rate falls. To examine how well this explanation fits the observed data, Figure 4.2 also plots the predicted contact hazard rate based solely on the number of new and old stocks on the other side of the market. The size of the fall in the contact rate between months 1 and 2 is quite well explained purely by the change in the size of the stocks, suggesting that the negative duration dependence is at least partially driven by the stock-flow mechanism.

However, other factors may also cause this pattern of duration dependence. For example, old agents might be less attractive to potential matching partners (old and new alike), simply because the fact that they remain unmatched for a long period signals inferior quality or high recruitment standards. Perhaps more importantly, selection bias can causes a decline in the contact hazards: if the contact hazards are heterogeneous for different individuals, those with high hazard rates exit faster, and as search duration lengthens, the slow-matching individuals would dominate the population.

Fig. 4.3 Contacts per month of search by unemployment spell length



One way to address this concern is by summarizing the contact patterns of different groups of individuals. We split the sample into four groups by the lengths of search spell: those who searched for less than or equal to 3 months, 3 to 6 months, 6 to 12 months, and longer than 12 months. The number of contacts over duration for each groups are summarized in Figure 4.3. The patterns suggest that the negative duration dependence in the full sample is not a pure artifact of selection. The fast matching group have more contacts than other groups especially in the first month, but all the groups on average make more contacts in the first month of search and have much lower contact rates in the second month onwards.[11]

4.5.2 Who contacts whom?

The stock-flow matching model has very clear predictions about who contacts whom: a job-seeker who have passed their screening period will never contact a vacancy in the old stock, and vice versa for vacancies. Our data record the identities of the job-seeker and vacancy pair which meet, allowing us to test about these predictions.

Recall that with a window period of $k = 4$ weeks, we classify agents with elapsed duration of ≤ 4 weeks as new, and > 4 weeks as old. And hence we count the contacts which fall into the four types by whether the agents are new or old. The numbers are displayed in Table 4.2, together with the average contact rates, calculated by the number of contacts over the number of months "at risk" for that type of agent.

Old agents have much lower contact rates than new agents, but when they do make contact, it is more likely to be with new agents on the other side of the market. Job-seekers on average have $c^U_{12} = 0.14$ contacts to old vacancies monthly when they are new, but this drops to $c^U_{22} = 0.06$ when they become old. Similarly, new vacancies receive $c^V_{21} = 0.78$ contacts per month from old job-seekers, and only $c^V_{22} = 0.20$ contacts from

[11] Even for the slowest-matching group, the difference in the contact rates between the first and second month is significant with $p < 0.001$ according to a paired t-test.

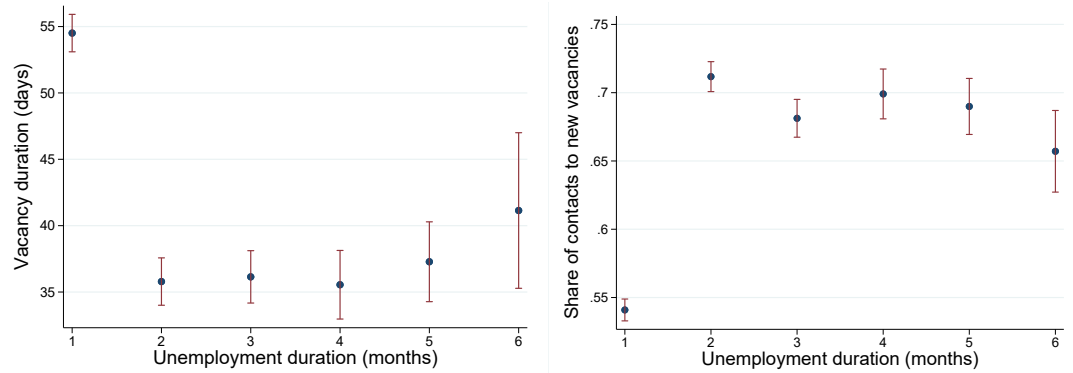
Table 4.2 Who contacts whom? Total contacts between new and old agents. The average contact rate is the total contacts divided by the number of job-seeker/vacancy months at risk. The average contact rates for the sample of active agents are in the square brackets.

	New job-seekers	Old job-seekers	Vacancy months at risk	
			All	[Active-only]
New vacancies	C_{11}=7,723	C_{21}=13,895	17,621	[8,683]
	$c^U_{11} = 0.17$ [0.52] $c^V_{11} = 0.44$ [0.89]	$c^U_{21} = 0.12$ [0.27] $c^V_{21} = 0.79$ [1.60]		
Old vacancies	C_{12}=6,093	C_{22} =6,934	35,259	[9,204]
	$c^U_{12} = 0.14$ [0.41] $c^V_{12} = 0.17$ [0.66]	$c^U_{22} = 0.06$ [0.13] $c^V_{22} = 0.20$ [0.75]		
Job-seeker months at risk	All 44,513 [Active-only] [14,894]	118,692 [52,275]		

old job-seekers when they become old. If we exclude the job-seekers and vacancies which did not have any contacts throughout the search spell, we observe similar changes in the contact rates: for job-seekers, the contacts to old vacancies fall from $c^U_{12} = 0.41$ when they are new to $c^U_{22} = 0.13$ when they are old, and for vacancies the contacts to old job-seekers fall from $c^V_{21} = 1.60$ to $c^V_{22} = 0.75$. However, Table 4.2 also shows that 20% of contacts happen between old job-seekers and old vacancies, something which is ruled out by pure stock-flow matching. Some of these may be accounted for by the changes in the preferences of job-seekers during searching. Although we do not fully observe such changes, our data do record the occupation and location preferences for job-seekers; roughly 75% of job-seekers did not change their preferences over the unemployment spell. If we exclude the job-seekers who have changed preferences, the fraction of stock-stock contacts fall to 17%, and the drop in the rate of contacts to old vacancies when a job-seeker becomes old is slightly larger (see Table 4.A1 in Appendix 4.A).

Figure 4.4 provides additional evidence that job-seekers are increasingly inclined to contact new vacancies as they stay in the market for longer. Panel (a) shows that, as

Fig. 4.4 Features of contacted vacancies over unemployment duration; 95% confidence intervals are given by the brackets.



(a) Mean duration of vacancies

(b) Share of contacts to new vacancies

Notes: We focus on the first 6 months of job search, for which we have more information: only 13% of unemployment spells lasted for over 6 months and only 7% of all contacts are conducted after the first 6 months of search.

job-seekers' search spells lengthen, the average duration of the vacancies they contact falls, and panel (b) shows that the share of applications to new vacancies increases with the job-seeker's search duration. Vacancies are also less likely to contact old job-seekers when they become old (as show by Figure 4.A1 in Appendix 4.A, which compares the distributions of the durations of contacted partners for new and old agents). These changes in the characteristics of contacted partners over search duration suggest that the type of friction specified by stock-flow matching is present in the data.

4.6 Econometric Evidence

4.6.1 Random matching model

We first estimate the random matching model for both job-seekers and vacancies. Table 4.3 reports estimates and model statistics for the full sample in the first two columns and for the active sample in third and forth column.

Table 4.3 Random matching, pooling old and new agents

	All searchers		Active searchers	
	Job-seekers	Vacancies	Job-seekers	Vacancies
$\log U$	-0.123^{***}	0.178^{***}	-0.131^{***}	0.170^{***}
	(0.008)	(0.033)	(0.008)	(0.035)
$\log V$	0.289^{***}	-0.122	0.084^{***}	-0.125^{*}
	(0.020)	(0.074)	(0.014)	(0.053)
Log Likelihood	-90,929	-46,643	-70,506	-36,656
Obs.	163,205	52,880	67,169	17,887

For job-seekers, the implied estimates of $\hat{\alpha} = 0.877$ (standard error 0.008) (since $\hat{\alpha} - 1 = -0.123$) and $\hat{\beta} = 0.289$ (0.020) are consistent with a contact function with constant returns to scale. For vacancies, the results also imply CRS, but the individual estimates of α and β are quite different, with $\hat{\alpha} = 0.178$ (0.033) and $\hat{\beta} = 0.878$ (0.074) (since $\hat{\beta} - 1 = -0.122$). The estimated congestion effects (i.e. $\alpha - 1$ for job-seekers and $\beta - 1$ for vacancies) tend to be small, and even insignificant for the vacancies. Since these estimates come from the same labour market, different estimates of α and β from the two sides of the market suggest some kind of model misspecification. When we restrict attention to the active job-seekers and vacancies, the results are very similar, except that impact of the vacancy stocks on contact rate of job-seeker is estimated to be smaller but is of similar significance level.

Since the stock-flow matching model predicts that the determinants of the contact rates change as the agent stay longer in the market, we also estimate the random matching model separately for old and new agents, as reported in Table 4.4. Panel (a) displays the results for all agents. When job-seekers are new, the total stock of vacancies has a strong positive effect on their contact rates. The effect of total vacancy stock shrinks significantly when the job-seeker becomes old, and the level of significance also drops. A more dramatic change of the same direction is observed when the active sample of job-seekers is used for the estimates, given in Panel (b). Analogous results are observed

among vacancies: the total stock of unemployed job-seekers has a strong impact on the contact rate of the new vacancies, yet it appears to have no predictive power over the contact rate of old vacancies. These patterns are not consistent with the random matching model.

Table 4.4 Random matching, separating old and new

(a) All searchers				
	Job-seekers		Vacancies	
	New	Old	New	Old
$\log U$	0.051^{***}	-0.075^{***}	0.375^{***}	0.093
	(0.010)	(0.017)	(0.022)	(0.051)
$\log V$	0.786^{***}	0.079^{*}	-0.077	-0.431^{**}
	(0.024)	(0.037)	(0.057)	(0.132)
Log Likelihood	-32,515	-54,964	-24,154	-18,768
Obs.	44,513	118,692	17,621	35,259
(b) Active searchers				
	Job-seekers		Vacancies	
	New	Old	New	Old
$\log U$	-0.071^{***}	-0.030	0.282^{***}	0.147^{**}
	(0.007)	(0.020)	(0.028)	(0.045)
$\log V$	0.539^{***}	0.058	-0.043	-0.361^{***}
	(0.014)	(0.035)	(0.044)	(0.098)
Log Likelihood	-20,938	-44,029	-17,782	-15,012
Obs.	14,894	52,275	8,683	9,204

4.6.2 Stock-flow matching model

Estimates of Equation (4.9) are reported in Table 4.5. The first column displays the estimates of $c^U_{1.}$ and $c^U_{2.}$, i.e. the contact rates for new and old job-seekers. The second column displays the estimates of $c^V_{.1}$ and $c^V_{.2}$, i.e. the contact rates for new and old vacancies. Column (1) shows that for old job-seekers the estimated $\hat{\pi}_2 = 0.226$ (standard error 0.039). Plugging this and the average stocks into Equation (4.7), we get an a-ratio (a_{22}/a_{21}) of 0.458, which suggests that an old job-seeker is estimated to be only 46% as

likely to contact a old vacancy as a new vacancy. A very similar effect is found among active job-seekers, with $\hat{\pi}_2 = 0.207(0.036)$ and $a_{22}/a_{21} = 0.482$ (column 3).

The stock-flow effects are also observed for vacancies. Column (2) suggests that for old vacancies, with $\hat{\pi}_4 = 0.170$ (0.054), the a-ratio is estimated to be 0.606, which means an old vacancy is only 61% as likely to be contacted by an old job-seeker compared to a new job-seeker. When only the active agents in the market are considered, column (4) gives $\hat{\pi}_4 = 0.203$ (0.050), and the a-ratio for old vacancies is estimated to be 0.506.

With these findings we reject the random matching prediction that the stock of new agents on the other side of the market has no additional effect on the contact rate.

It is also interesting to examine the determinants of contact rates of the new agents, although the stock-flow theory does not provide any clear predictions. For both new job-seekers and new vacancies, the aggregate stocks on the other side of the market has a strong positive impact on their contact rates. Also when we account for the stock-flow effects by incorporating the effects of the new stocks, the estimated parameters of the matching function are similar for old and new agents. For instance, we have $\hat{\alpha} = 0.965$ (0.010) and $\hat{\beta} = 0.307$ (0.029) for the new job-seekers, and $\hat{\alpha} = 1.027$ (0.016) and $\hat{\beta} = 0.234$ (0.027) for the old job-seekers. The differences are small, but are statistically significant with $p < 0.001$ for $\hat{\alpha}$ and $p = 0.001$ for $\hat{\beta}$.

We also estimate the conditional fixed-effects model, using which we fully control for both observed and unobserved individual heterogeneity and relax the assumption that individual characteristics are uncorrelated with the independent variables. However, the use of the conditional fixed-effects model results in a far more restrictive sample, since agents who are observed for only one month are excluded. The estimates are reported in Table 4.B1 in Appendix 4.B.

The results from conditional fixed-effects estimates are very similar to those of the random-effects estimate. In column (1), we see that the stock of new vacancies has no

Table 4.5 Random effect Poisson model for monthly contacts. Table reports estimates of the contact rate for new and old job-seekers and vacancies as a function of stocks of new and old agents, as in (4.9).

	(a) All searchers		*(b) Active searchers*	
	Job-seekers	Vacancies	Job-seekers	Vacancies
New				
$\rho_1(\log V)$	0.573^{***}	-0.042	0.350^{***}	0.076
	(0.058)	(0.135)	(0.040)	(0.066)
$\rho_2(\log v)$	-0.266^{***}	-0.273^{**}	-0.033	-0.272^{***}
	(0.063)	(0.085)	(0.040)	(0.064)
$\rho_3(\log U)$	0.452^{***}	0.336^{***}	0.237^{***}	0.260^{***}
	(0.028)	(0.057)	(0.026)	(0.049)
$\rho_4(\log u)$	-0.487^{***}	-0.047	-0.307^{***}	0.017
	(0.028)	(0.058)	(0.026)	(0.055)
Old				
$\pi_1(\log V)$	0.008	-0.093	-0.060	0.043
	(0.042)	(0.155)	(0.043)	(0.083)
$\pi_2(\log v)$	0.226^{***}	-0.221^{*}	0.207^{***}	-0.317^{***}
	(0.039)	(0.097)	(0.036)	(0.082)
$\pi_3(\log U)$	0.156^{***}	-0.010	0.020	-0.005
	(0.030)	(0.081)	(0.036)	(0.065)
$\pi_4(\log u)$	-0.130^{***}	0.170^{**}	-0.017	0.203^{***}
	(0.024)	(0.054)	(0.025)	(0.050)
New				
α	0.965	0.290	0.930	0.277
	(0.010)	(0.027)	(0.008)	(0.028)
β	0.307	0.684	0.316	0.804
	(0.029)	(0.099)	(0.018)	(0.051)
$\alpha+\beta$	1.272	0.974	1.246	1.081
	(0.034)	(0.112)	(0.018)	(0.054)
Old				
α	1.027	0.160	1.003	0.198
	(0.016)	(0.049)	(0.020)	(0.042)
β	0.234	0.685	0.147	0.726
	(0.027)	(0.102)	(0.027)	(0.062)
$\alpha+\beta$	1.261	0.845	1.150	0.924
	(0.026)	(0.105)	(0.017)	(0.053)
a-ratio	0.458	0.606	0.482	0.506
Log Likelihood	-87,502	-43,917	-65,766	-33,833
Obs.	163,205	52,880	67,169	17,887

additional effect on the contact rate of new job-seekers: $\hat{\rho}_2 = -0.069$ (0.099), which is not significantly different from zero. But the stock of new vacancies has a significant and positive effect on the contact rate of old job-seekers: $\hat{\pi}_2 = 0.288$ (0.048). Similar patterns are observed on the vacancy side of the market. The stock of new job-seekers has a significant and positive effect on the contact rate of old vacancies, but has no effect on new vacancies. Comparing the a-ratios between the random- and fixed-effects Poisson models, we see that the ratios are smaller, i.e. even further from the random matching assumption, in the fixed-effects estimates. The similarity of the random- and fixed-effects models is consistent with the hypothesis that aggregate stocks of U and V in each labour market are largely uncorrelated with unobserved individual effects in the contact rate.

4.6.3 Robustness

Our results thus far are based on the assumption that the window period is approximately one month. In other words, agents can arrange interviews with potentially suitable partners within one month. We have argued that given the database and screening inherent in the Matching Service matching system, this is a reasonable assumption. But it is worth exploring whether our results change if we change the window period. Studies on employer search have suggested that vacancies are filled within short periods of time (e.g. Barron et al., 1999; Burdett and Cunningham, 1998; Coles and Smith, 1998), which implies that employers may require a shorter initial search period. This is also supported by the average lengths of search spells in our sample: the job-seeker spends 62 days searching, whereas it takes the median employer only 25 days to fill a vacancy. Also, employers tend to have more experience with the Matching Service, as many of them use the system repeatedly for various vacancies, compared to job-seekers who typically only one search spell. Andrews et al. (2013) therefore use $k^U = 4$ weeks as the window

period for job-seekers and $k^V = 2$ weeks for vacancies, and checked with several other window periods, which turn out to give similar results.

We also estimate the model of contact rates using other combinations of window periods, including k^U-k^V of 2-2, 3-2, 4-2, 3-3, 4-3, 5-4, 5-5, 6-5, and 6-6 weeks. Results are reported in Table 4.B2 in Appendix 4.B. The stock-flow effects (i.e. $\hat{\pi}_2$ for job-seekers and $\hat{\pi}_4$ for vacancies) remain positive and significant for all combinations but the one with $k^V = 6$ weeks. In general the stock-flow effects in contacts are not very sensitive to the choice of window periods.

Our data contains repeated search spells for around 25% of job-seekers and half of vacancies. The search activities of job-seekers who have earlier search experience might differ from those who are in their first spell. And for the vacancy panel, the repeated spells are places in multiple-place vacancy orders. Therefore as a robustness check, we estimate Equation (4.9) using the first search spells observed in the data and report the results in Table 4.6. The estimated stock-flow effects are virtually unchanged when the model for those who search for the first time. For old job-seekers, $\hat{\pi}_2 = 0.190$ (0.045), and for old vacancies $\hat{\pi}_2 = 0.215$ (0.050).[12]

Another concern is that the apparent stock-flow effects shown in Table 4.5 arise from aggregation across distinct sub-markets with different labour market tightness. To see this, suppose that there are two distinct sub-markets, one with excess vacancies and one with excess job-seekers. Matching in both sub-markets is random, but job-seekers in the first market match quickly, while vacancies in the second market match quickly. If we aggregate the two sub-markets, we observe new job-seekers matching old vacancies (sub-market 1) and new vacancies matching old job-seekers (sub-market 2).

[12] There might be a reason to expect the stock-flow effects to be stronger: Kuo and Smith (2009) find that, among job-seekers who have repeated unemployment spells, the new job-seekers who are only shortly apart from their last spell are less likely to match with the old vacancy stock (and behave like an old job-seeker), because they have screened those vacancies during their previous search. If repeated unemployment spells are separated by short employment spells, excluding the later spells would strengthen the stock-flow effects. However, such an effect is not present in our data.

Table 4.6 Random effect Poisson model for monthly contacts, first spells only.

	All searchers		Active Searchers	
	Job-seeker	Vacancies	Job-seeker	Vacancies
New				
$\rho_1(\log V)$	0.777^{***}	0.196^{*}	0.493^{***}	0.131
	(0.068)	(0.087)	(0.053)	(0.080)
$\rho_2(\log v)$	-0.550^{***}	-0.469^{***}	-0.235^{***}	-0.296^{***}
	(0.075)	(0.069)	(0.059)	(0.060)
$\rho_3(\log U)$	0.410^{***}	0.369^{***}	0.271^{***}	0.248^{***}
	(0.034)	(0.048)	(0.032)	(0.059)
$\rho_4(\log u)$	-0.445^{***}	-0.044	-0.315^{***}	0.020
	(0.034)	(0.051)	(0.032)	(0.050)
Old				
$\pi_1(\log V)$	-0.003	0.087	-0.046	0.119
	(0.055)	(0.094)	(0.043)	(0.098)
$\pi_2(\log v)$	0.190^{***}	-0.306^{***}	0.170^{***}	-0.310^{***}
	(0.045)	(0.091)	(0.043)	(0.076)
$\pi_3(\log U)$	0.180^{***}	-0.047	0.048	-0.142
	(0.036)	(0.057)	(0.032)	(0.078)
$\pi_4(\log u)$	-0.150^{***}	0.215^{***}	-0.018	0.284^{***}
	(0.023)	(0.050)	(0.025)	(0.055)
a-ratio	0.516	0.549	0.543	0.398
Log Likelihood	-62,996	-34,089	-46,108	-29,173
Obs.	120,563	24,955	46,059	15,842

In practice, sub-markets could be separated by skill-levels, occupations or geographic locations. The most straightforward to define are the geographical sub-markets. The Matching Service have offices in different local districts, which can be grouped into the Central, West and East labour markets. The vast majority of contacts occur within each sub-market. We calculate the series of old and new stocks for each market, and estimate Equation (4.9) for each of the sub-markets. The results are reported in Table 4.7.

The West labour market is unambiguously the largest among the three in terms of both numbers of job-seekers and vacancies. The estimated stock-flow effect is significant for this market only, and the size of such effects are similar to those estimated for the overall labour market. For the Central and East labour markets, the estimates suggest that only the total stocks of agents on the other side of the market are effective determinants

Table 4.7 Random effect Poisson model for monthly contacts, separate estimates for three geographic sub-markets

	(a) Central		*(b) West*		*(c) East*	
	Job-seekers	Vacancies	Job-seekers	Vacancies	Job-seekers	Vacancies
New						
$\rho_1(\log V)$	0.780^{***}	-0.146	0.554^{***}	-0.095	0.634^{***}	-0.131
	(0.073)	(0.187)	(0.050)	(0.129)	(0.100)	(0.104)
$\rho_2(\log v)$	-0.247^{**}	-0.356^{*}	-0.445^{***}	-0.020	-0.188	-0.206^{**}
	(0.081)	(0.159)	(0.057)	(0.095)	(0.096)	(0.076)
$\rho_3(\log U)$	0.497^{***}	0.246^{**}	0.390^{***}	0.341^{***}	0.447^{***}	0.491^{***}
	(0.052)	(0.093)	(0.034)	(0.076)	(0.049)	(0.061)
$\rho_4(\log u)$	-0.544^{***}	0.002	-0.373^{***}	-0.072	-0.572^{***}	-0.094
	(0.050)	(0.081)	(0.036)	(0.070)	(0.052)	(0.056)
Old						
$\pi_1(\log V)$	0.402^{***}	-0.507^{**}	-0.045	0.110	0.336^{***}	-0.552^{***}
	(0.063)	(0.156)	(0.038)	(0.159)	(0.091)	(0.164)
$\pi_2(\log v)$	-0.009	-0.188	0.122^{***}	-0.202^{*}	0.004	0.307^{*}
	(0.053)	(0.204)	(0.034)	(0.094)	(0.071)	(0.152)
$\pi_3(\log U)$	0.118	0.411^{***}	0.260^{***}	-0.152	-0.143^{**}	0.154
	(0.070)	(0.104)	(0.036)	(0.092)	(0.050)	(0.099)
$\pi_4(\log u)$	-0.112^{*}	-0.183	-0.205^{***}	0.213^{***}	0.053	0.078
	(0.048)	(0.097)	(0.022)	(0.062)	(0.044)	(0.077)
Log Likelihood	-21,858	-10,281	-42,840	-23,923	-21,785	-9,356
Obs.	39,117	12,705	66,327	30,910	57,761	9,265

of the contact rates. However, we notice that for all three labour markets, the estimated impact of total vacancy (unemployment) stocks on the contact rate of new job-seekers (vacancies) is always positive and significant, whereas the impact on old job-seekers (vacancies) is insignificant and close to zero. This result is more plausible under the stock-flow matching model compared to the random matching model. Therefore, although the stock-flow effects appear to be weakened by disaggregation, we argue that this could result from a smaller sample size which makes it harder to identify the effects.

4.7 Conclusions

In a two-sided search environment where both sides of the market have access to a database of potential partners, it seems natural to consider the idea that search frictions (and hence periods of search) are generated by the lack of suitable partners, rather than

a lack of information about those partners.[13] Since internet job-search has become so widespread, studying such an environment seems increasingly relevant. Although our dataset pre-dates internet job search, it comes from a two-sided database, and provides a unique opportunity to study the mechanism by which job-seekers and vacancies contact each other.

Our findings suggest that the stock-flow model outperforms the traditional random matching model in explaining the observed pattern of contacts in our data. Two key features emerge. First, contact rates exhibit strong negative duration dependence. We show that is unlikely to be due to the sample selection effect, because both fast-exiting and slow-exiting agents make more contacts in their first month of search than in the rest of their search spell. The extent of this duration dependence is quite well captured by the fact that the stock of available partners drops once the stock has been sampled. Second, job-seekers are more likely to contact old vacancies when they are new, and less so when they become old (the average exposure of vacancies contacted over search duration decreases and the fraction of application to new vacancies increases), and analogous patterns are observed for vacancies.

We also directly estimate the stock-flow effects on the contact rates of job-seekers and vacancies, distinguishing the old and new agents. The stock-flow effects appear to be strong and significant, suggesting that stock-flow contacts are much more likely to happen than stock-stock contacts. Such a finding is robust to different definitions of old and new agents and the sample selection, and the exclusion of inactive agents in the construction of variables. However, disaggregation the labour market into sub-markets segregated by geographical boarders does appear to weaken the stock-flow effects.

[13] This chapter does not reconcile the estimates of α and β from the two sides of the market. As demonstrated in the Chapter 3, the fact that estimates differ widely between the job-seeker model and the vacancy model suggests there is some degree of misspecification.

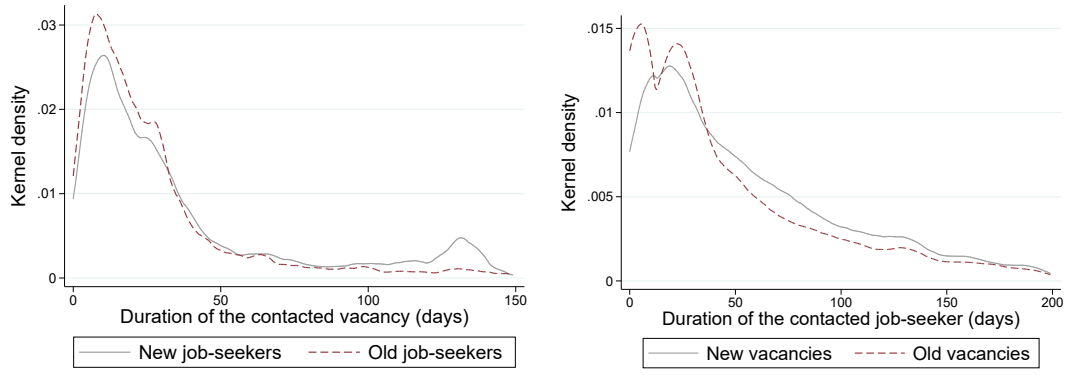
Our results confirm those of Andrews et al. (2013) and a number of other empirical papers, but crucially we are able to show that the mechanism arises via the contact technology rather than the match probability. Since the argument of the stock-flow theory relates to contacts rather than matches, this provides the most direct evidence so far that stock-flow matching describes labour market search behaviour better than models of random matching.

Appendix 4.A Additional descriptive statistics

Table 4.A1 Who contacts whom? This table repeats the analysis of Table 4.2, but excludes the job-seekers whose occupational choices change during their search spell.

	New job-seekers	Old job-seekers
New vacancies	6,435 $c_{11}^U = 0.17$	8,642 $c_{21}^U = 0.10$
Old vacancies	5,339 $c_{12}^U = 0.14$	4,255 $c_{22}^U = 0.04$
Job-seeker months at risk	37,878	86,667

Fig. 4.A1 Old agents tend to contact those with a shorter duration compared to new agents, as suggested by the distributions



(a) Durations of vacancies contacted by new and old job-seekers

(b) Durations of job-seekers contacted by new and old vacancies

Notes: Panel (a) shows that job-seekers are more likely to contact older vacancies when they are new compared to when they are old; panel (b) shows the same duration pattern for vacancies; the Kolmogorov-Smirnov tests reject the equality of distributions for the duration of contacted agents in both panels with $p < 0.001$.

Appendix 4.B Additional regression results

Table 4.B1 Conditional fixed-effects model for monthly contacts, which can be compared to the random-effects estimates in Table 4.5.

	(a) All searchers		*(b) Active searchers*	
	Job-seekers	Vacancies	Job-seekers	Vacancies
New				
$\rho_1(\log V)$	0.119	−0.226	−0.137	−0.351
	(0.116)	(0.308)	(0.096)	(0.260)
$\rho_2(\log v)$	−0.069	−0.038	−0.044	0.086
	(0.099)	(0.164)	(0.081)	(0.144)
$\rho_3(\log U)$	0.436***	0.122	0.489***	0.178
	(0.055)	(0.113)	(0.047)	(0.131)
$\rho_4(\log u)$	−0.351***	0.012	−0.358***	0.024
	(0.055)	(0.088)	(0.047)	(0.100)
Old				
$\pi_1(\log V)$	−0.171**	−0.075	−0.292***	−0.063
	(0.058)	(0.212)	(0.060)	(0.202)
$\pi_2(\log v)$	0.288***	−0.134	0.288***	−0.166
	(0.048)	(0.119)	(0.046)	(0.127)
$\pi_3(\log U)$	0.109*	−0.345**	0.036	−0.300**
	(0.044)	(0.112)	(0.047)	(0.100)
$\pi_4(\log u)$	−0.134***	0.319***	−0.126***	0.311***
	(0.027)	(0.078)	(0.031)	(0.071)
New				
α	1.085	0.135	1.130	0.202
	(0.022)	(0.077)	(0.030)	(0.096)
β	0.050	0.736	−0.181	0.735
	(0.062)	(0.217)	(0.082)	(0.196)
$\alpha+\beta$	1.135	0.870	0.950	0.937
	(0.052)	(0.192)	(0.064)	(0.144)
Old				
α	0.975	−0.027	0.910	0.011
	(0.027)	(0.064)	(0.032)	(0.055)
β	0.117	0.791	−0.004	0.771
	(0.048)	(0.184)	(0.046)	(0.147)
$\alpha+\beta$	1.092	0.765	0.906	0.782
	(0.052)	(0.195)	(0.057)	(0.141)
a-ratio	0.376	0.402	0.374	0.363
Obs.	65,027	13,936	65,027	13,936
Log Likelihood	-34,916	-13,509	-34,908	-13,504

Table 4.B2 Random effect Poisson model for monthly contact using various window periods

(a) Job-seeker panel									
k^U-k^V:	2-2	3-2	4-2	3-3	4-3	5-4	5-5	6-5	6-6
New									
$\rho_1(\log V)$	0.448***	0.502***	0.612***	0.514***	0.584***	0.480***	0.474***	0.367***	0.341***
	(0.067)	(0.068)	(0.041)	(0.061)	(0.068)	(0.054)	(0.041)	(0.051)	(0.045)
$\rho_2(\log v)$	−0.071	−0.247***	−0.319***	−0.261***	−0.274***	−0.152*	−0.143**	−0.052	−0.019
	(0.072)	(0.068)	(0.047)	(0.068)	(0.066)	(0.061)	(0.053)	(0.062)	(0.052)
$\rho_3(\log U)$	0.446***	0.459***	0.427***	0.461***	0.441***	0.432***	0.436***	0.475***	0.484***
	(0.038)	(0.041)	(0.025)	(0.035)	(0.035)	(0.028)	(0.030)	(0.033)	(0.031)
$\rho_4(\log u)$	−0.623***	−0.551***	−0.463***	−0.556***	−0.477***	−0.471***	−0.476***	−0.532***	−0.541***
	(0.033)	(0.038)	(0.026)	(0.033)	(0.034)	(0.028)	(0.027)	(0.030)	(0.029)
Old									
$\pi_1(\log V)$	0.057	0.023	0.006	0.057	0.028	−0.065	−0.112*	−0.180***	−0.208***
	(0.036)	(0.032)	(0.035)	(0.045)	(0.044)	(0.050)	(0.043)	(0.053)	(0.054)
$\pi_2(\log v)$	0.244***	0.260***	0.267***	0.186***	0.211***	0.243***	0.292***	0.322***	0.357***
	(0.026)	(0.024)	(0.032)	(0.035)	(0.040)	(0.047)	(0.037)	(0.056)	(0.055)
$\pi_3(\log U)$	0.025	0.046*	0.174***	0.036	0.175***	0.242***	0.246***	0.293***	0.285***
	(0.018)	(0.021)	(0.031)	(0.021)	(0.027)	(0.032)	(0.026)	(0.029)	(0.033)
$\pi_4(\log u)$	−0.141***	−0.109***	−0.146***	−0.098***	−0.144***	−0.179***	−0.183***	−0.230***	−0.225***
	(0.019)	(0.021)	(0.024)	(0.016)	(0.019)	(0.023)	(0.022)	(0.027)	(0.025)
Log Likelihood	-87789	-87686	-87447	-87733	-87501	-87635	-87630	-88067	-88061
Obs.	163205	163205	163205	163205	163205	163205	163205	163205	163205

(b) Vacancy panel									
k^U-k^V:	2-2	3-2	4-2	3-3	4-3	5-4	5-5	6-5	6-6
New									
$\rho_1(\log V)$	−0.102	−0.110	−0.117	0.023	0.051	−0.068	0.007	0.020	0.159
	(0.111)	(0.112)	(0.103)	(0.122)	(0.115)	(0.114)	(0.132)	(0.146)	(0.105)
$\rho_2(\log v)$	−0.261**	−0.287**	−0.289**	−0.335***	−0.378***	−0.254***	−0.303**	−0.316**	−0.375***
	(0.099)	(0.104)	(0.105)	(0.090)	(0.093)	(0.077)	(0.103)	(0.112)	(0.084)
$\rho_3(\log U)$	0.389***	0.333***	0.315***	0.302***	0.231**	0.372***	0.342***	0.328***	0.247***
	(0.075)	(0.090)	(0.076)	(0.062)	(0.074)	(0.074)	(0.057)	(0.082)	(0.068)
$\rho_4(\log u)$	0.046	0.106	0.118	0.019	0.091	−0.079	−0.070	−0.056	−0.022
	(0.095)	(0.110)	(0.091)	(0.075)	(0.073)	(0.081)	(0.064)	(0.084)	(0.062)
Old									
$\pi_1(\log V)$	−0.138	−0.169	−0.181*	−0.071	−0.091	−0.121	−0.120	−0.128	−0.147
	(0.105)	(0.098)	(0.086)	(0.109)	(0.107)	(0.135)	(0.168)	(0.169)	(0.165)
$\pi_2(\log v)$	−0.007	−0.008	−0.004	−0.202*	−0.196*	−0.202*	−0.209	−0.202	−0.309*
	(0.068)	(0.068)	(0.076)	(0.078)	(0.078)	(0.100)	(0.120)	(0.114)	(0.129)
$\pi_3(\log U)$	0.037	0.059	0.051	0.010	0.018	0.024	0.051	0.040	0.221*
	(0.046)	(0.053)	(0.044)	(0.067)	(0.079)	(0.079)	(0.075)	(0.079)	(0.105)
$\pi_4(\log u)$	0.121**	0.085*	0.087*	0.162**	0.143*	0.136*	0.120*	0.130*	0.063
	(0.042)	(0.035)	(0.035)	(0.052)	(0.064)	(0.054)	(0.061)	(0.063)	(0.090)
Log Likelihood	-44580	-44586	-44585	-44151	-44155	-43925	-43967	-43969	-44291
Obs.	52880	52880	52880	52880	52880	52880	52880	52880	52880

Chapter 5

Conclusions

In this thesis, we have reported three studies which investigate different research questions on the topic of labour market search. In one of the studies (Chapter 2) we explore social comparison effects in job search decisions using a laboratory experiment. In two other studies (Chapters 3 and 4) we use unique observational data. In Chapter 3 we decompose the matching process into the meeting process and matching probability. Chapter 4 tests predictions of the stock-flow matching model in the meeting process.

In Chapter 2, we find that within a search task there is a reference wage effect: when a partner exits, the subject chooses a new reservation wage which is increasing in partner income. And we also observe a social learning effect: between search tasks, subjects who have been paired with a more patient and successful partner increase their reservation wages in the next task. These social comparison effects can contribute to the correlation in the search outcomes in the peer groups observed in the real world.

We are aware that there are at least two concerns regarding our interpretations of the findings, which can potentially be addressed by additional treatments. The first concern is that the observed effects could be the combined effects of observing others (which facilitates social comparisons) and being *observed by* others, which is a mechanism

overlooked in the current study. One could disentangle the effects of being observed and observing others using a treatment in which one of the paired subjects receives search feedback from the partner whereas their own search feedback remains private. Another concern is that the reference wage effects in our experiment might be attributed to some complicated anchoring effects, in the sense that subjects are responding to the income and offer information they receive *without* considering it as decisions made and outcomes received by another person. Although this is not a particularly compelling explanation for various systematic effects we observe, to fully rule out this possibility, one can test another deviation from our social comparison treatments, by paying the paired subjects for different search sessions, such that the partner's decision has no payoff consequences. One would expect the the reference wage effects to be absent in this new treatment, if they are caused by distributional preferences.

We also note that in reality social interactions are much richer than the minimum interactions that we introduce in the laboratory setting (i.e. the exchange of search feedback with an anonymous other). Therefore an interesting direction for further research is to investigate the link between laboratory choice and real job-search outcomes. The link could be at aggregate level, e.g. Brown et al. (2011) relate the findings of declining reservation wage in the lab to the observational studies which also report a declining reservation wage. Perhaps the link at individual level is even more interesting, as it fully addresses the external validity concerns not only for the observed effects but also for the underlying mechanisms, e.g. DellaVigna and Paserman (2005) show that the impatience elicited in the laboratory tasks can be used to predict real search outcomes. In the context of social comparison effects one might test the linkage between the individual job-seeker's distributional preferences elicited in the lab and the real search outcomes, controlling measures of employment status and search outcomes in the individual's peer groups.

In Chapter 3, we decompose the matching process into the meeting process and the matching probability conditional on a meeting for both sides of the labour market. To achieve consistency between the job-seeker and vacancy estimates, we use a specification which accounts for temporal aggregation bias and geographic spill-overs, according to the suggestions of two separate strands of the empirical literature. This exercise gives rise to three key findings. First, the aggregate stocks of job-seekers and vacancies mainly affect the matching rate of the individual agents through affecting their meeting rate rather than their matching probability. Second, the decline of the matching hazard is mainly driven by the decline of the meeting hazard as the agents spend longer in the market. These two findings apply to both job-seekers and vacancies. And third, various characteristics affect the matching rate through affecting both the meeting rate and the matching probability.

In Chapter 4, we test the implications of stock-flow matching in the patterns of interviews, and report various evidence which favors stock-flow matching relative to random matching: the inflow rate of new agents, rather than the total stock of agents, determines the contact rates of existing agents. This contributes to the empirical literature on the evidence of stock-flow matching, because the test on interviews is a more direct test about the search mechanism, compared to the previous tests which focus on matches. The findings may also serve as a potential input for future theoretical analysis.

Important policy implications can be drawn from the findings of Chapter 3 and Chapter 4. Chapter 3 suggests that cutting unemployment benefits to increase the matching probability for the unemployed might not be very effective, because the decline of the matching hazard arises from a substantial decline in the meeting hazard. Also, many unemployment benefits are contingent on active job search efforts. However, the findings from Chapter 4 suggest that this might fail to promote re-employment and reduce long-term unemployment: the reason job-seekers make fewer contacts as they

search for longer is that there are not any potential partners which they could meet and match with.

We recognise that objections to these implications may arise from the fact that our decomposition is not perfectly mapped to the simple theoretical framework of stochastic search described in Pissarides (2000), in which the meeting rate is determined by market size and tightness, and the matching probability is determined by reservation wage and reservation profit. Specifically, for example, the meeting decisions we observe may partially reflect reservation wages and search efforts of individual job-seekers.

The data employed by Chapter 3 and Chapter 4 are quite unique in the sense that they document events in a well-defined labour market and record a rich set of agent-level information. Another interesting question for further research is to test the stability of the matching function and the meeting function, and the robustness of the stock-flow patterns using similar datasets of a more recent period or at a larger scale. This will address the concern that our findings might result from the particular matching technology in our sample, as contacts between job-seekers and vacancies are moderated by the Matching Service, which might have certain policy for arranging the interviews.

And last but not least, we note that the richness of our data facilitates the analysis without requiring structural assumptions, yet it also affords plenty of opportunities to "folk" a specific result (e.g. using different specifications), which casts doubt on the credibility of any analysis. In keeping with most of the literature, we attempt to alleviate this problem by reporting several specifications but there might be a more effective approach. Future research might consider machine learning methods (e.g. Guerrero and López, 2015, for an application of the methods on estimating the matching function) to exploit even more details in information recorded in the data, without making choices over specifications.

References

Abbring, J., van den Berg, G., Gautier, P., Lomwel, A., van Ours, J. and Ruhm, C. (2002), Displaced workers in the United States and the Netherlands, *in* P. Kuhn, ed., 'Losing Work, Moving on: International Perspectives on Worker Displacement', W.E. Upjohn Institute.

Abeler, J., Altmann, S., Kube, S. and Wibral, M. (2010), 'Gift exchange and workers' fairness concerns: When equality is unfair', *Journal of the European Economic Association* **8**(6), 1299–1324.

Akerlof, G. A. and Yellen, J. L. (1990), 'The fair wage-effort hypothesis and unemployment', *Quarterly Journal of Economics* **105**(2), 255–283.

Altmann, S., Armin, F., Jäger, S. and Zimmermann, F. (2016), Learning about job search: A field experiment with job seekers in Germany. IZA Discussion Paper No. 9040.

Alvarez de Toledo, P., Nunez, F. and Usabiaga, C. (2011), 'An empirical analysis of the matching process in andalusian public employment agencies', *Hacienda Publica Espanola* **3**, 67–102.

Andrews, M. J., Bradley, S., Stott, D. and Upward, R. (2008), 'Successful employer search? An empirical analysis of vacancy duration using micro data', *Economica* **75**, 455–480.

Andrews, M. J., Bradley, S., Stott, D. and Upward, R. (2013), 'Estimating the stock-flow matching model using micro data', *Journal of the European Economic Association* **11(5)**, 1153–1177.

Andrews, M. J., Bradley, S. and Upward, R. (2001), 'Estimating the probability of a match using microeconomic data for the youth labour market', *Labour Economics* **8**, 335–357.

Angelova, V., Güth, W. and Kocher, M. G. (2012), 'Co-employment of permanently and temporarily employed agents', *Labour Economics* **19**(1), 48–58.

Arulampalam, W. (2001), 'Is unemployment really scarring? effects of unemployment experience on wages', *Economic Journal* **111**(475), F585–F606.

Bandura, A. and Walters, R. H. (1977), *Social learning theory*, Prentice-Hall Englewood Cliffs, NJ.

Barron, J., Berger, M. and Black, D. (1999), 'Employer search, training and vacancy duration', *Economic Inquiry* **35**, 167–192.

Belot, M., Kircher, P. and Muller, P. (2016), Providing advice to job seekers at low cost: An experimental study on on-line advice. IZA Discussion Paper No. 10068.

Berman, E. (1997), 'Help wanted, job needed: estimates of a matching function from Employment Service data', *Journal of Labor Economics* **15**, S251–92.

Bertrand, M. and Duflo, E. (2017), 'Field experiments on discrimination', *Handbook of Economic Field Experiments* **1**, 309–393.

Bicchieri, C. and Xiao, E. (2009), 'Do the right thing: but only if others do so', *Journal of Behavioral Decision Making* **22**(2), 191–208.

Blanchard, O. J. and Diamond, P. A. (1994), 'Ranking, unemployment duration and wages', *Review of Economic Studies* **61**, 417 – 434.

Blau, D. M. and Robins, P. K. (1986), 'Job search, wage offers, and unemployment insurance', *Journal of Public Economics* **29**(2), 173 – 197.

Bloemen, H. G. and Stancanelli, E. G. (2001), 'Individual wealth, reservation wages, and transitions into employment', *Journal of Labor Economics* **19**(2), 400–439.

Bohnet, I. and Zeckhauser, R. (2004), 'Social comparisons in ultimatum bargaining', *Scandinavian Journal of Economics* **106**(3), 495–510.

Bramoullé, Y. and Saint-Paul, G. (2010), 'Social networks and labor market transitions', *Labour Economics* **17**(1), 188–195.

Braunstein, Y. M. and Schotter, A. (1982), 'Labour market search: an experimental study', *Economic Inquiry* **20**(1), 133–144.

Brown, M., Flinn, C. J. and Schotter, A. (2011), 'Real time search in the laboratory and the market', *American Economic Review* **101**(2), 948–974.

Brown, M., Setren, E. and Topa, G. (2016), 'Do informal referrals lead to better matches? evidence from a firm's employee referral system', *Journal of Labor Economics* **34**(1), 161–209.

Burda, M. C. and Profit, S. (1996), 'Matching across space: evidence on mobility in the czech republic', *Labour Economics* **3**(3), 255–278.

Burdett, K. and Coles, M. G. (1999), 'Long-term partnership formation: Marriage and employment', *Economic Journal* **109**(456), pp. F307–F334.

Burdett, K., Coles, M. and Van Ours, J. (1994), Temporal aggregation bias in stock-flow models. Discussion Paper No. 967, CEPR, London.

Burdett, K. and Cunningham, E. J. (1998), 'Towards a theory of vacancies', *Journal of Labour Economics* **16**, 445–478.

Burgess, S. and Profit, S. (2001), 'Externalities in the matching of workers and firms in britain', *Labour Economics* **8**(3), 313–333.

Bursztyn, L., Ederer, F., Ferman, B. and Yuchtman, N. (2014), 'Understanding mechanisms underlying peer effects: Evidence from a field experiment on financial decision', *Econometrica* **82**, 1273–1301.

Cai, H., Chen, Y. and Fang, H. (2009), 'Observational learning: Evidence from a randomized natural field experiment', *American Economic Review* **99**(3), 864–82.

Cai, J., De Janvry, A. and Sadoulet, E. (2015), 'Social networks and the decision to insure', *American Economic Journal: Applied Economics* **7**(2), 81–108.

Calvó-Armengol, A. and Jackson, M. O. (2004), 'The effects of social networks on employment and inequality', *American Economic Review* **94**(3), 426–454.

Cappellari, L. and Tatsiramos, K. (2015), 'With a little help from my friends? Quality of social networks, job finding and job match quality', *European Economic Review* **78**, 55–75.

Card, D., Mas, A., Moretti, E. and Saez, E. (2012), 'Inequality at work: The effect of peer salaries on job satisfaction', *American Economic Review* **102**(6), 2981–3003.

Charness, G. and Kuhn, P. (2007), 'Does pay inequality affect worker effort? experimental evidence', *Journal of Labor Economics* **25**(4), 693–723.

Cingano, F. and Rosolia, A. (2012), 'People I know: Job search and social networks', *Journal of Labor Economics* **30**(2), 291–332.

Clark, A. E. (2003), 'Unemployment as a social norm: Psychological evidence from panel data', *Journal of Labor Economics* **21**(2), 323–351.

Clark, A. E., Masclet, D. and Villeval, M. C. (2010), 'Effort and comparison income: Experimental and survey evidence', *ILR Review* **63**(3), 407–426.

Clark, A. E. and Oswald, A. J. (1996), 'Satisfaction and comparison income', *Journal of Public Economics* **61**(3), 359 – 381.

Cobb-Clark, D. A., Frijters, P. and Kalb, G. R. (2004), Do you need a job to find a job? IZA Discussion Paper 1211.

Cohn, A., Fehr, E., Herrmann, B. and Schneider, F. (2014), 'Social comparison and effort provision: Evidence from a field experiment', *Journal of the European Economic Association* **12**(4), 877–898.

Coles, M. G. and Muthoo, A. (1998), 'Strategic bargaining and competitive bidding in a dynamic market equilibrium', *Review of Economic Studies* **65(2)**, 235–260.

Coles, M. G. and Smith, E. (1998), 'Marketplaces and matching', *International Economic Review* **39(1)**, 239–254.

Coles, M., Jones, P. and Smith, E. (2010), 'A picture of stock-flow unemployment in the United Kingdom', *Macroeconomic dynamics* **14**, 427–453.

Coles, M. and Petrongolo, B. (2008), 'A test between stock-flow matching and the random matching function approach', *International Economic Review* **49**, 1113–1141.

Coles, M. and Smith, E. (1996), 'Cross-section estimation of the matching function: evidence from England and Wales', *Economica* **63**, 589–97.

Conley, T. G. and Udry, C. R. (2010), 'Learning about a new technology: Pineapple in Ghana', *American Economic Review* **100**(1), 35–69.

Cox, J. and Oaxaca, R. (2000), 'Good news and bad news: Search from unknown wage offer distributions', *Experimental Economics* **2**, 197–225.

Dal Bó, P. and Fréchette, G. R. (2017), 'On the determinants of cooperation in infinitely repeated games: A survey', *Journal of Economic Literature* **forthcoming**.

Davis, S. J., Faberman, R. J. and Haltiwanger, J. C. (2013), 'The establishment-level behavior of vacancies and hiring', *Quarterly Journal of Economics* **581**, 622.

DellaVigna, S. and Paserman, M. D. (2005), 'Job search and impatience', *Journal of Labor Economics* **23**(3), 527–588.

Diamond, P. A. (1982), 'Aggregate demand management in search equilibrium', *Journal of Political Economics* **90**, 881–894.

Dustmann, C., Glitz, A. and Schönberg, U. (2016), 'Referral-based job search networks', *Review of Economic Studies* **83**(2), 514–546.

Ebrahimy, E. and Shimer, R. (2010), 'Stock-flow matching', *Journal of Economic Theory* **145**, 1325–1353.

Eckstein, Z. and van den Berg, G. (2007), 'Empirical labour search: a survey', *Journal of Econometrics* **136**, 531–564.

Erens, B. and Hedges, B. (1990), *Survey of Incomes in and out of Work*, HMSO.

Faberman, R. J. and Kudlyak, M. (2014), The intensity of job search and search duration. The Federal Reserve Bank of Richmond Working Paper.

Falk, A., Huffman, D. and Sunde, U. (2009*a*), Do I have what it takes? Equilibrium search with type uncertainty and non-participation. IZA Discussion Paper No. 2531.

Falk, A., Huffman, D. and Sunde, U. (2009*b*), Self-confidence and search. IZA Discussion Paper No. 2525.

Fehr, E. and Schmidt, K. M. (1999), 'A theory of fairness, competition, and cooperation', *Quarterly Journal of Economics* **114**, 817–868.

Festinger, L. (1954), 'A theory of social comparison processes', *Human Relations* **7**(2), 117–140.

Fischbacher, U. (2007), 'z-Tree: Zurich toolbox for ready-made economic experiments', *Experimental Economics* **10**(2), 171–178.

Frank, R. H. (1984), 'Are workers paid their marginal products?', *American Economic Review* **74**(4), 549–571.

Frank, R. H. (1985), *Choosing the right pond: Human behavior and the quest for status.*, New York: Oxford University Press.

Friedl, A., De Miranda, K. L. and Schmidt, U. (2014), 'Insurance demand and social comparison: An experimental analysis', *Journal of Risk and Uncertainty* **48**(2), 97–109.

Frijters, P. and van der Klaauw, B. (2006), 'Job search with nonparticipation', *Economic Journal* **116**(508), 45–83.

Gächter, S., Nosenzo, D. and Sefton, M. (2012), 'The impact of social comparisons on reciprocity', *Scandinavian Journal of Economics* **114**(4), 1346–1367.

Gächter, S., Nosenzo, D. and Sefton, M. (2013), 'Peer effects in pro-social behavior: Social norms or social preferences?', *Journal of the European Economic Association* **11**(3), 548–573.

Galenianos, M. (2013), 'Learning about match quality and the use of referrals', *Review of Economic Dynamics* **16**(4), 668–690.

Gamba, A., Manzoni, E. and Stanca, L. (2017), 'Social comparison and risk taking behavior', *Theory and Decision* **82**(2), 221–248.

García-Pérez, J. I. (2006), 'Job separation in a non-stationary search model: a structural estimation to evaluate alternative unemployment insurance systems', *Journal of Applied Econometrics* **21**(2), 245–272.

Gregg, P. and Petrongolo, B. (2005), 'Stock-flow matching and the performance of the labor market', *European Economic Review* **49**, 1987–2011.

Greiner, B. (2015), 'Subject pool recruitment procedures: organizing experiments with ORSEE', *Journal of the Economic Science Association* **1**(1), 114–125.

Guerrero, O. A. and López, E. (2015), 'Firm-to-firm labor flows and the aggregate matching function: A network-based test using employer–employee matched records', *Economics Letters* **136**, 9–12.

Haisley, E., Mostafa, R. and Loewenstein, G. (2008), 'Subjective relative income and lottery ticket purchases', *Journal of Behavioral Decision Making* **21**(3), 283–295.

Harrison, G. W. and Morgan, P. (1990), 'Search intensity in experiments', *Economic Journal* **100**, 478–486.

Hey, J. D. (1981), 'Are optimal search rules reasonable? And vice versa?(And does it matter anyway?)', *Journal of Economic Behavior & Organization* **2**(1), 47–70.

Hey, J. D. (1982), 'Search for rules for search', *Journal of Economic Behavior & Organization* **3**(1), 65–81.

Hey, J. D. (1987), 'Still searching', *Journal of Economic Behavior & Organization* **8**(1), 137–144.

Horváth, G. (2014), 'Occupational mismatch and social networks', *Journal of Economic Behavior & Organization* **106**, 442–468.

Ibanez, M., Czermak, S. and Sutter, M. (2009), 'Searching for a better deal–on the influence of group decision making, time pressure and gender on search behavior', *Journal of Economic Psychology* **30**(1), 1–10.

Ioannides, Y. M. and Loury, L. D. (2004), 'Job information networks, neighborhood effects, and inequality', *Journal of Economic Literature* **42**(4), 1056–1093.

Jarosch, G. and Pilossoph, L. (2015), Statistical discrimination and duration dependence in the job finding rate, Technical report, working paper, Princeton University.

Jensen, P. and Westergård-Nielsen, N. C. (1987), 'A search model applied to the transition from education to work', *Review of Economic Studies* **54**(3), 461–472.

Kettemann, A., Mueller, A. I. and Zweimüller, J. (2017), Wages, workers and vacancy durations: Evidence from linked data. Mimeo, University of Zurich.

Knez, M. J. and Camerer, C. F. (1995), 'Outside options and social comparison in three-player ultimatum game experiments', *Games and Economic Behavior* **10**(1), 65–94.

Kogut, C. A. (1990), 'Consumer search behavior and sunk costs', *Journal of Economic Behavior & Organization* **14**(3), 381–392.

Kroft, K., Lange, F. and Notowidigdo, M. J. (2013), 'Duration dependence and labor market conditions: Evidence from a field experiment', *Quarterly Journal of Economics* **128**, 1123–1167.

Krueger, A. B. and Mueller, A. (2011), 'Job search, emotional well-being and job finding in a period of mass unemployment: Evidence from high-frequency longitudinal data', *Brookings Papers on Economic Activity* **42**(1 (Spring), 1–81.

Krupka, E. and Weber, R. A. (2009), 'The focusing and informational effects of norms on pro-social behavior', *Journal of Economic Psychology* **30**(3), 307–320.

Kuhn, P. and Mansour, H. (2014), 'Is internet job search still ineffective?', *Economic Journal* **124**(581), 1213–1233.

Kuhn, P. and Skuterud, M. (2004), 'Internet job search and unemployment durations', *American Economic Review* **94**(1), 218–232.

Kuo, M.-Y. and Smith, E. (2009), 'Marketplace matching in britain: Evidence from individual unemployment spells', *Labour Economics* **16**, 37–46.

Lahno, A. M. and Serra-Garcia, M. (2015), 'Peer effects in risk taking: Envy or conformity?', *Journal of Risk and Uncertainty* **50**(1), 73–95.

Linde, J. and Sonnemans, J. (2012), 'Social comparison and risky choices', *Journal of Risk and Uncertainty* **44**, 45–72.

Linde, J. and Sonnemans, J. (2015), 'Decisions under risk in a social and individual context: The limits of social preferences?', *Journal of Behavioral and Experimental Economics* **56**, 62–71.

Loewenstein, G. F., Thompson, L. and Bazerman, M. H. (1989), 'Social utility and decision making in interpersonal contexts', *Journal of Personality and Social psychology* **57**(3), 426.

Lollivier, S. and Rioux, L. (2010), 'An empirical examination of the sources of changes over time in the job finding rate using reservation wages and rejected wage offers', *International Economic Review* **51**(4), 1039–1069.

Machin, S. and Manning, A. (1999), The causes and consequences of long-term unemployment in Europe, *in* D. Card and O. Ashenfelter, eds, 'Handbook of Labor Economics', North-Holland, pp. 3085–3139.

Manski, C. F. (1993), 'Identification of endogenous social effects: The reflection problem', *Review of Economic Studies* **60**, 531–542.

Manski, C. F. (2000), 'Economic analysis of social interactions', *Journal of Economic Perspectives* **14**, 115–136.

Marmaros, D. and Sacerdote, B. (2002), 'Peer and social networks in job search', *European Economic Review* **46**, 870–879.

McCall, J. J. (1970), 'Economics of information and job search', *Quarterly Journal of Economics* **84**(1), 113–126.

McGee, A. and McGee, P. (2016), 'Search, effort, and locus of control', *Journal of Economic Behavior & Organization* **126**, 89 – 101.

Moon, P. and Martin, A. (1990), 'Better heuristics for economic search: Experimental and simulation evidence', *Journal of Behavioral Decision Making* **3**(3), 175–193.

Moon, P. and Martin, A. (1992), 'Purchasing decisions, partial knowledge, and economic search experimental and simulation evidence', *Journal of Behavioral Decision Making* **5**(4), 253–266.

Moon, P. and Martin, A. (1996), 'The search for consistency in economic search', *Journal of Economic Behavior & Organization* **29**, 311–321.

Oaxaca, R. L. and Cox, J. C. (1989), 'Laboratory experiments with a finite horizon job search model', *Journal of Risk and Uncertainty* **2**, 301–329.

Oaxaca, R. L. and Cox, J. C. (1992), 'Direct test of the reservation wage property', *Economic Journal* **102**, 1423–1432.

Offerman, T. and Schotter, A. (2009), 'Imitation and luck: An experimental study on social sampling', *Games and Economic Behavior* **65**(2), 461–502.

Patacchini, E. and Zenou, Y. (2007), 'Spatial dependence in local unemployment rates', *Journal of Economic Geography* **7**(2), 169–191.

Petrongolo, B. and Manning, A. (2017), 'How local are labor markets? evidence from a spatial job search model', *American Economic Review* **forthcoming**.

Petrongolo, B. and Pissarides, C. (2006), 'Scale effects in markets with search', *Economic Journal* **116**(January), 21–44.

Petrongolo, B. and Pissarides, C. A. (2001), 'Looking into the black box: A survey of the matching function', *Journal of Economic Literature* **39(2)**, 390–431.

Pissarides, C. (1979), 'Job matchings with state employment agencies and random search', *Economic Journal* **89**, 818 – 833.

Pissarides, C. (1990), *Equilibrium Unemployment Theory*, Basil Blackwell.

Pissarides, C. (2000), *Equilibrium Unemployment Theory*, 2nd edition edn, Basil Blackwell.

Pissarides, C., Layard, R. and Hellwig, M. (1986), 'Unemployment and vacancies in Britain', *Economic Policy* **1**, 499–559.

Riach, P. and Rich, J. (2006), 'An experimental investigation of sexual discrimination in hiring in the English labor market', *The B.E. Journal of Economic Analysis & Policy* **5**(2), 1–22.

Rohde, I. M. T. and Rohde, K. I. M. (2011), 'Risk attitudes in a social context', *Journal of Risk and Uncertainty* **43**(3), 205–225.

Rosholm, M. (2000), 'Observed and unobserved heterogeneity in the duration dependency parameter', *Economics Letters* **66**(3), 347–351.

Russo, G. and van Ommeren, J. (1998), 'Gender differences in recruitment outcomes', *Bulletin of Economic Research* **50**, 155–66.

Sasaki, M. (2008), 'Matching function for the Japanese labour market: Random or stock-flow?', *Bulletin of Economic Research* **60**, 209–230.

Schlag, K. H. (2011), Imitation and social learning, Technical report, Encyclopaedia of the Sciences of Learning (forthcoming).

Schmitt, J. and Wadsworth, J. (1993), 'Unemployment benefit levels and search activity', *Oxford Bulletin of Economics and Statistics* **55**(1), 124.

Schotter, A. and Braunstein, Y. M. (1981), 'Economic search: an experimental study', *Economic Inquiry* **19**, 1–25.

Schunk, D. (2009), 'Behavioral heterogeneity in dynamic search situations: Theory and experimental evidence', *Journal of Economic Dynamics & Control* **33**, 1719–1738.

Schunk, D. and Winter, J. (2009), 'The relationship between risk attitudes and heuristics in search tasks: A laboratory experiment', *Journal of Economic Behavior & Organization* **71**(2), 347–360.

Schwerter, F. (2016), Social reference points and risk taking, Technical report, Bonn Econ Discussion Papers, University of Bonn.

Shimer, R. (2007), 'Mismatch', *American Economic Review* **97**, 1074–1101.

Shimer, R. (2012), 'Reassessing the ins and outs of unemployment', *Review of Economic Dynamics* **15**(2), 127–148.

Soetevent, A. R. and Bruzikas, T. (2016), Risk and loss aversion, price uncertainty and the implications for consumer search. Tinbergen Institute Discussion Paper 16-049/VII.

Soininen, H. (2007), Stock-flow matching and directed search - Finnish evidence. Working Paper.

Sonnemans, J. (1998), 'Strategies of search', *Journal of Economic Behavior & Organization* **35**, 309–332.

Sonnemans, J. (2000), 'Decisions and strategies in a sequential search experiment', *Journal of Economic Psychology* **21**(1), 91–102.

Stepner, M. (2013), 'BINSCATTER: Stata module to generate binned scatterplots', Statistical Software Components, Boston College Department of Economics.

Stern, S. (1989), 'Estimating a simultaneous search model', *Journal of Labor Economics* **7**(3), 348–369.

Teyssière, G. (1995), 'Matching processes in the labour market: an econometric study', *Labour Economics* **2**(4), 421–435.

van den Berg, G. (1990), 'Nonstationarity in job search theory', *Review of Economic Studies* **57**(2), 255–277.

van Ours, J. and Ridder, G. (1991), 'Cyclical variations in vacancy durations and vacancy flows: an empirical analysis', *European Economic Review* **35**, 1143–1155.

van Ours, J. and Ridder, G. (1992), 'Vacancies and the recruitment of new employees', *Journal of Labor Economics* **10**(2), 138–155.

Wooldridge, J. M. (2010), *Econometric Analysis of Cross Section and Panel Data*, MIT Press.

Zajonc, R. B. (1965), 'Social facilitation', *Science* **149**(3681), 269–274.