Signals in two-sided search

Friedrich Gerd Poeschel
Brasenose College
University of Oxford

A thesis submitted in partial fulfilment
of the requirements for the degree of
Doctor of Philosophy

Trinity Term 2011

Word count: 64000 words

Gewidmet denen,
die auch bei
einem Scheitern dieser Arbeit
immer gern
von mir gehört hätten

Abstract

Signals in two-sided search

Friedrich Gerd Poeschel, Brasenose College

D.Phil. Thesis, Trinity Term 2011

We introduce signals to search models of two-sided matching markets and explore the implications for efficiency. In a labour market model in which firms can advertise wages and workers can choose effort, we find that advertisements can help overcome the Diamond paradox. Advertisements fix workers' beliefs, so that workers will react if firms renege on advertisements. Firms then prefer to advertise truthfully. Next, we consider a market with two-sided heterogeneity in which types are only privately observable. We identify a simple condition on the match output function for agents to signal their types truthfully and for the matching to exhibit positive assortative matching despite search frictions. While our theoretical work implies that the efficiency of matching increases as information technology spreads, empirical matching functions typically suggest that it declines. By estimating more general matching functions, we show that the result of declining efficiency can partly be attributed to omitted variable bias.

JEL Classification Numbers: C78, D83, J31, J63, J64

Keywords: search, matching, signals, Diamond paradox, sorting, matching function

Acknowledgements

I would like to thank my supervisors Professor Godfrey Keller and Professor Fabien Postel-Vinay for very helpful discussions and suggestions. Chapters 2 and 3 were written under the patient supervision of Godfrey Keller, while chapter 4 was written under that of Fabien Postel-Vinay. My work on chapter 2 has benefited greatly from conversations with Jim Albrecht, Leo Kaas, James Malcomson, Andrew Rhodes, Emily Troscianko, and Randy Wright. I have also received useful comments from audiences at the University of Oxford, Nuffield College, RWI Essen, the University of Vienna, the labour economics workshop at the University of Mainz, and at Ecole Polytechnique.

For very useful comments and suggestions on chapter 3, I am grateful to Hiroyuki Adachi, Siddhartha Bandyopadhyay, Melvyn Coles, Bruno Decreuse, Jan Eeckhout, Pieter Gautier, Leo Kaas, Philipp Kircher, Franco Peracchi, Christopher Pissarides, Lones Smith, Margaret Stevens and to seminar participants at the Royal Economic Society Easter School 2008, the Econometric Society World Congress 2010, the XXth Aix-Marseille Doctoral Spring School and at the universities of Hannover, Oxford, Rome "Tor Vergata", Konstanz, Hamburg, Munich and at the Ecole Polytechnique. This chapter was partly written while I was a visitor at the Ecole normale supérieure, and I thank this institution for its hospitality.

Finally, chapter 4 has benefited from conversations with Robert Solow and Gesine Stephan. Most of this chapter was written at the Ecole Polytechnique, and I thank that institution for its hospitality. I am grateful to internet cafés in Palaiseau and Paris for giving me access to Stata during this time. Financial support from the Economic and Social Research Council (№ PTA-031-2004-00250), from the Centre Cournot pour la Recherche en Economie, and from the EU's 6th Framework Programme is gratefully acknowledged. All remaining errors are my own.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Motivation and focus

1.1.1 Signals in search

Economic agents search, both in many real-world markets and in the models of such markets: workers search for jobs, firms search for employees, consumers search for the best deal. Let us consider search behaviour in very general terms. An agent X searches in order to respond to some need. After considering T possible solutions one after the other, where T is chosen by X, she selects one of these solutions or gives up.

This simple search process can lead to a variety of outcomes. Typically, however, X considers $T > 1$ possible solutions. Then there remains at least one possible solution that was considered but not selected, and considering it may have cost resources. Next, if a solution is selected, this will typically not be the best solution that X could have found. Rather, X will often content herself with a solution that is sufficiently good so that investing in further search appears unjustified. Yet it might also happen that an agent undertakes unjustified investments in further search (or any search at all) after misjudging the probability of finding a better solution (or any acceptable solution at all).

Under perfect information, none of these cases could arise: then X will engage in search if and only if it makes sense for her, and she will consider only the solution she most prefers among all feasible solutions. After all, she would not knowingly spend resources on considering solutions she will not select. Hence, she would select the first solution she considers. When she can identify the best solution that is feasible for her, she would optimally choose to consider this best solution first and would never content herself with anything less than this best solution.

Typical outcomes of search behaviour are therefore inefficient when compared to a full-information benchmark, and this inefficiency by comparison arises precisely from the lack of information.[1] Since the root cause underlying the inefficiencies of search is imperfect information, the inefficiencies may be addressed by improvements in the available information. In particular, X would like to know the answers to the following questions:

1. Where do I find the best feasible solution?
2. Does the solution I can expect to end up with justify further search (or any search)?

Armed with such information, optimising agents would search *selectively*: they would only consider solutions that are worth considering and would only consider the best among these. Some degree of selective search does not only result from efficiency considerations on a theoretical level but also appears much more plausible as a description of actual search processes in the real world than "blind" sequential search. Imagine to the contrary that real-world agents could not search selectively at all. Then they might sometimes invest resources in (further) search when such investments are not justified. Hence agents would be willing to pay up to the expected value of such investments in return for precise information on when (further) search is justified. Similarly, an agent's search will not terminate after considering the first solution, unless all solutions are known to be the same or the agent strikes it lucky with the first solution. An agent would thus be willing to pay up to the search costs incurred from considering solutions that are not selected in return for information that allows to avoid considering such solutions.

[1] That search is only necessitated by a lack of information has been emphasised by contributors as early as McCall (1970).

The existence of such strong incentives to make information available implies that the situation assumed by "blind" sequential search would be unstable in practice. One would expect that agents seek and find a way to search more selectively. Moreover, one may observe that agents in the real world do search selectively rather than randomly, which suggests that they have somehow obtained advance information. This thesis analyses how signalling allows agents to search more selectively. Adopting the terminology introduced by Spence (1973), we define signalling as strategic communication through deliberately chosen messages, or signals. Such signals can transmit information between agents. In particular, signals can provide the answers to the two questions above. Then agents need no longer gather information through search. In effect, agents can replace costly search among potential solutions by essentially costless search among signals thereof.

By making search obsolete, signals can remove the inefficiencies associated with search. In the ideal case, a situation of perfect information arises endogenously as a consequence of information transmission through signals. However, signals will assume this function only if they are *truthful*. Because signals are strategically chosen, they do not necessarily have any informational content. Worse, they allow agents to deliberately mislead other agents whenever there is an incentive to lie. Hence the credibility of signals will be a key concern in this thesis.

1.1.2 Two-sided matching markets

To explore and analyse how signals can make search efficient, one would want to consider a decentralised market where agents search, both in the real world and in theoretical models. The labour market and various markets for consumer goods appear as leading examples for such markets. Yet the former and the latter differ so fundamentally that one cannot necessarily analyse them in the same way. In fact, the literature now treats labour markets as belonging to a different class of markets that may be labelled *matching markets*. The marriage market, the market for rented housing, and the market for venture capital are all often seen to fall into the same class (see Adachi (2003), for example). We next briefly characterise matching markets before we explain the difference between two-sided markets of this class and two-sided markets for consumer goods.

A matching market features economic agents who make offers to trade. For matching markets to clear, it is necessary to join (*to match*) single offers to trade into durable coalitions (*matches*). As a result of market clearing, some overall allocation (*a matching*) of the offers emerges. While each agent may well be behind several offers, for simplicity one can treat each offer to trade as independent of other offers by the same agent. Then one can abstract from the distinction between an agent and her offer to trade, as if the market matched agents and not their offers. Dropping this distinction also emphasises a crucial characteristic of matching markets: agents whose offers are matched form an economic partnership, and an offer's market value therefore typically depends on the agent making the offer.

To illustrate, a job seeker's offer of work is matched with a vacancy in the labour market, leading to an employment relation between the job seeker and the firm. In the marriage market, two agents' offers of companionship are matched, leading to a marriage. The value of these offers depends on the associated agents: job seekers might well prefer an otherwise identical vacancy in a more productive or prestigious firm. Similarly, the companionship of a famous actor with, unsurprisingly, movie-star looks might be considered more desirable than the companionship of a dull accountant. Finally, we call the matching market *two-sided* whenever the matches are pairs and the two matched offers of each pair are always drawn from the same two sets of offers to trade, one from each. Following this simple definition, a market for consumer goods with only buyers and sellers who transact in pairs is also two-sided.[2]

We now identify three important differences between two-sided matching markets and two-sided markets for consumer goods. As a useful example, consider the market for rented housing (a matching market) and the market for houses as a consumer

[2] Roth and Sotomayor (1990) define a two-sided market somewhat differently as one where each agent is confined to one side. They consider labour markets two-sided, but not a market for consumer goods, as "any particular agent might be a buyer at one price and a seller at another" (p. 1). Yet in the labour market, too, any particular agent might buy in labour at one price and sell her own labour at another. Likewise, agents in the market for rented housing may switch from letting properties to renting themselves in reaction to price changes. It is admittedly less easy to imagine such changes happening in the marriage market.

good. Firstly, the matches formed in a matching market are in principle *durable* and thereby lead to a prolonged partnership between agents. Thus a tenancy agreement may be regarded as repeated interaction over time between the same two agents. By contrast, the sale of a house is a one-off transaction, and the associated agents might never meet again. Secondly, the behaviour of an agent during a partnership is often not observable ex ante, but affects the value of the offer. This makes the offer a so-called "experience good" whose quality can only be verified by trying it out for a while. Consider a student who rents a room in a property with a resident landlady. How good an idea this is will only become clear to the student after living there for a while. Perhaps the student discovers that the landlady habitually plays the trumpet to salute the break of day. By contrast, any such habits on the side of the seller of a house do not affect the quality of living in that house and therefore leave the value of the house unaffected. Thirdly, and again due to the fact that the offer's value depends on an agent, offers in a matching market are *not tradable* in the sense that they cannot be sold on. While the buyer of a house is entirely free to sell it on, the aforementioned student acquires the right to use the room but may not sublet it to any person the landlady would rather not have in her house.

Indeed, Diamond's (1982) model of search on a market for consumer goods lets agents interact only once; they immediately consume the acquired good (which is thus never traded twice) and separate again to return to a stage of production. However, models of search on a matching market often do not explicitly use any of the distinctive features of a matching market and therefore carry over to markets for consumer goods despite the differences we emphasised. For this reason, also some material in this thesis (particularly chapter 3) can be applied to both classes of markets.

Yet we preferred to write this thesis with a focus throughout on the class of matching markets, which we would justify as follows. For matching markets and above all labour markets, search now appears to be the dominant theoretical approach (see Rogerson et al. (2005)). By contrast, the search-theoretic approach to markets for consumer goods appears to face much fiercer competition from alternative approaches. That apart, it seemed to us that signals in search are more central to real-world behaviour on matching markets, in the sense that search models without signals seem less realistic here. For example, random search on the market for houses implies that each and every property agent encountered is visited, which might conceivably describe some agents' real-world behaviour. Meanwhile, random search on the marriage market implies that every potential partner encountered is dated, which is hard to imagine as real-world behaviour.

1.1.3 Signals in real-world matching

Consider a developed labour market in the real world. In our view, both job seekers and firms make ample use of signals for their search, with a beneficial effect on the efficiency of their search. When an employer advertises a vacancy, the information given about job and employer constitutes a signal of the job offer's quality. A job seeker may observe the advertisement and become interested in this offer. In this case (and only in this case), the job seeker sends an application to the employer, and the information therein constitutes a signal of the job seeker's offer. If the employer likewise becomes interested in this particular job seeker an interview will be arranged. Hence, real-world agents are in a position to condition their willingness to meet on observing a certain signal.

In the interview, both sides have the opportunity to test the information they have thus far received about the offers. If and only if both accept the trade, the vacancy is filled, and neither the vacancy nor the job seeker's offer to work are on the market any longer. A very similar process can just as well be initiated the other way around: a job seeker might advertise himself on an internet platform, or might just send applications to a variety of employers he could imagine working for. An employer who is interested in the job seeker's offer will respond by giving details of the job that would be offered to the job seeker. If this job attracts the job seeker's interest, an interview will be arranged, and so forth. We would point out that observing signals in these procedures is associated with next to no costs, while interviews create travel costs, opportunity costs, or preparation costs.

A large and growing share of the activity in two-sided matching markets now happens on online platforms. On these platforms, participating agents of both sides maintain profiles of themselves. These are designed to cover the most relevant pieces of information about the offer. However, it is the agent who chooses the information displayed in her profile, so that the information does not have to be truthful. The profiles can be viewed at practically no (marginal) cost, while both sides would incur certain costs from meeting up. Any agent can search the profiles and can show interest in a particular offer. The agent who makes this particular offer is then notified and will have a look at the offer made to her, using the profiles. Hence the profiles are the signals here, and as before mutual interest precedes the meeting in person. In fact, we would argue that the growing use of online platforms places ever more emphasis on the role of signals.[3]

Similar procedures are easily discovered in the housing market, too. Before any meeting is arranged, the landlord has typically advertised the room(s) and potential tenants have introduced themselves on the phone, so that the landlord has already information about the occupational status, the sex, the age, any smoking habits, and perhaps the previous living arrangements of the potential tenant. Likewise, a date on the marriage market is normally only arranged when both sides have become interested in one another on the grounds of information already obtained. This information may be signalled through looks, dress style, body language, and the like. It is telling that dates without prior information are called 'blind dates'.

1.2 Literature review

1.2.1 The search-theoretic approach

The role of two-sided matching markets is to pair up offers to trade from two sets. A large and growing literature has been devoted to the study of this problem on a theoretical level. Many contributions in the literature are sufficiently general or

[3] Autor (2001) suggests that the growth of such online platforms is due to the advantages, in terms of costs and otherwise, of online job search and recruitment over previous search channels. Yet he also finds that these technologies greatly exacerbate the problem of false signals because the cost difference between truthful and false signals has often disappeared.

adaptable to be relevant for more than one real-world example, although most models were built with a view towards the labour market (see Rogerson et al. (2005) for a survey), while few contributions consider the market for rented housing or the market for venture capital. We would broadly structure this literature by the kind of *search frictions* included in a model. Search frictions are obstacles or difficulties that have to be overcome before a match can be concluded. They are useful for classification because several kinds of search frictions are rarely present simultaneously without one kind being much more important than others. We distinguish four main kinds of search frictions: discounting, explicit search costs, congestion, and no frictions at all.

Discounting Discounting of future payoffs creates difficulties in matching whenever it takes some time (at least in expectation) before an agent meets another agent so that a match could possibly result. Then this time will be valuable in the sense that agents would like to reduce the time elapsing as far as possible, and any time elapsed is an indirect cost for agents. Therefore, the case of discounting is also referred to as indirect or implicit search costs. For example, contributions by Burdett and Coles (1997) and by Shimer and Smith (2000) feature discounting as search friction.

Explicit search costs Perhaps the simplest search frictions are strictly positive costs that an agent has to pay out of her pocket whenever she meets another agent. As opposed to indirect search costs, they are called direct or explicit search costs. Morgan (1998) and Atakan (2006) build search models with explicit search costs.

Congestion A third type of frictions arises when several agents simultaneously try to match with a particular agent, although only one of them can succeed. At the same time, it might be that nobody tries to match with another, yet identical agent. This can happen even when there are only two agents on each side if agents do not coordinate. Hence these frictions are also called coordination frictions. They feature in models by Burdett et al. (2001) and Shimer (2005a), for example.

No frictions Whenever agents can meet potential match partners without any difficulties, search is frictionless. Among others, Gale and Shapley (1962) and Becker (1973) let search be frictionless in their models of a matching market.

Models in the last category may be called *matching models* as opposed to models in the other categories, which may be called *search models*. Since the use of these termini in the literature is not always consistent, our classification will inevitably re-label some contributions.[4]

1.2.2 Search models

Following early work by Stigler (1961), a literature exemplified by the work of McCall (1970), Mortensen (1970), Burdett (1978), and Pissarides (2000) emerged around the idea of *sequential search*: each agent samples only one offer at a time whose quality is modelled as a random draw from the distribution of qualities in the market. Alternatively, the quality of any resulting match is made a random draw from some distribution. In either case, all offers appear homogeneous ex ante, i.e. before the draw.

Search models do not include frictions just to make life complicated. Instead, search frictions are included because they are a pervasive phenomenon in reality: in the vast majority of real-world search situations, agents cannot match without difficulties and expending resources on their search. Importantly, agents' optimal behaviour changes when there are search frictions, as follows.

As an agent proceeds with her search and samples one offer after the other, each time she faces the question whether to end the search by accepting the current offer, or whether to continue searching in the hope of finding a better deal. If the existence of search frictions makes search costly in some way, agents will rationally choose to settle for less than the best feasible match. To see this, suppose the search has produced the opportunity to conclude a match where the quality of the offer thus

[4] For example, the model by Burdett and Coles (1997) with discounting is here classified as a search model while labelled a matching model by its authors.

obtained is close to that of the best feasible match. Then the gain from any further search is at most the difference between this offer's quality and the quality of the offer in the best feasible match. Yet when trying to reap this gain by continuing to search, additional search costs are inevitably incurred, too. Therefore, if the quality of the current offer is close enough to that in the best feasible match, the expected gain from continued search will no longer outweigh the expected costs of continued search. How close is close enough? This is a question an agent needs to answer for herself by choosing some cut-off quality level. The agent's optimal policy then exhibits the *reservation wage property*: a potential match is accepted if and only if the offer's quality weakly exceeds her cut-off level.

While this problem of optimal stopping was initially analysed in static environments such as that in Weitzman (1979), since Diamond (1971) the attention of the literature has increasingly turned to equilibrium search models. In these models, both sides of the market are modelled as agents who behave strategically, so that the wage distribution faced by a job seeker, for example, is endogenous. Also the models in this thesis are equilibrium search models. For the empirical estimation of search models, Flinn (1986) and van den Berg and Ridder (1998) may be considered milestones.

1.2.3 Matching models

The classical set-up of matching models features two groups of offers, A and B, say. Each agent associated with an offer in group A ranks all offers in group B according to her preferences, and each agent associated with an offer in B likewise ranks all offers in A. Then some algorithm or a central matchmaker who behaves in a specified way is invoked to produce a matching. The widely-used algorithm by Gale and Shapley (1962) has agents from group A successively offering the trade to their most preferred potential match partners in group B who have not yet rejected them. The agents in group B retain or reject the offers received and finally only retain the most preferred offer of all. Alternatively, the agents in B successively make offers, while the agents in A stay put, but the resulting matching will in general differ.

Apart from Gale and Shapley (1962), Koopmans and Beckmann (1957) and Roth and Sotomayor (1990) exemplify this literature. It gives special attention to the properties of the resulting matching (as a function of the matching algorithm employed). One question of interest is whether a *stable matching* is produced, that is, no two agents would rather want to be matched to each other than to their current match partners, and no agent would prefer not to be matched at all. In other words, optimising agents change ex post any matching produced by some algorithm if the matching is not stable. Next, an algorithm may systematically favour one group over the other. For the algorithm by Gale and Shapley (1962), it can be shown that the group proposing the trade fares systematically better in the resulting matching than the group that stays put.

The matching models by Becker (1973) and Lam (1988) assume neither any algorithm nor a matchmaker, but leave the task of market clearing to decentralised interaction among optimising agents. They are therefore closer in spirit to search models. The properties of matchings then depend on agents' utility functions and in particular on the *match production function*. The specification of the match production function determines how the match partners' offers to trade, playing the role of inputs, jointly generate output during the match. It is then examined under which conditions the resulting matching exhibits sorting among the types or ranks of offers. In particular, *positive assortative matching* obtains when the type of the offer just acceptable to an agent (the marginal type) is increasing in the type of her own offer, so that there is a positive correlation between the types of matched offers. In our chapter 3, matching models and above all Becker (1973) will serve as a benchmark because one can think of a matching model as the limit point of a search model where search frictions become negligible (see Adachi (2003)).

1.2.4 Adding signals

Mortensen (1976) made a strong conceptual case for signals as one way among others to improve the matching process. However, while he calculates the potential welfare gain attainable with signals, he does not model signalling itself nor search processes when signals are available. The very general argument by Mortensen (1976) in favour of signals in search has received a rather specific answer in the literature: the vast majority of models introducing signals to search fall into the literature of *directed search* (or competitive search).[5]

Directed search models are typically cast in labour market terms, a choice we will follow throughout this section. In these models, firms transmit information by publicly advertising job characteristics so that workers can select the firm they apply to according to the advertised information, an approach used by Ramaswami (1983), Montgomery (1991), Peters (1991), Moen (1997), and others. The structure common to most directed search models can be represented as interaction in three stages between firms and homogeneous workers. Firstly, all firms simultaneously choose wages (or contracts, e.g. wage-tenure profiles) that they advertise and commit to them. Next, all workers observe the advertised wage offers and simultaneously choose which firm to apply to in this period. Finally, job seekers and firms match as follows. A firm that received exactly one application hires this one applicant. A firm that received several applications somehow selects, perhaps randomly, one of these applicants for the job. Firms that did not receive any applications can try again in the next period, as can workers who are not hired. Higher advertised wages tend to attract more applications. In their choice which wage to advertise, firms thus face a trade-off between attracting applicants and keeping wage costs low. Similarly, workers face a trade-off between jobs with high wages but low probability of being hired and jobs with low wages but high probability of being hired. The result of these models typically is that, in equilibrium, workers are indifferent as to which firm to apply to because the expected wage is the same across firms.

We would argue that there is a major problem with this entire literature: no justification is given why firms should be able to commit to their advertisements. Yet this would be necessary, since firms have an incentive to reduce their wage offer once workers have applied to them: as long as the wage reduction is not too large,

[5] Outside the directed search literature, there are only very few contributions that consider signals in search. Among these few, most contributions allow for neither discounting nor explicit search costs, but consider the frictionless case (see e.g. Hoppe et al. (2009) and Coles et al. (2011)).

workers will grudgingly still accept, as they would otherwise incur the search costs again. In other words, by the logic explained in section 1.2.2, a wage offer close enough to the advertised wage will also be accepted. Advertised contracts or mechanisms can likewise be made ex post less favourable for workers. Hence firms have a strong incentive to renege on their advertisements. Since the work of Diamond (1971), firms' ability to set wages ex post has been known as a fundamental problem that can lead to complete market breakdown even with infinitesimally small search frictions. Indeed, we show formally in chapter 2 that only the assumption of commitment keeps the market in a directed search model from breaking down, despite the ad-hoc nature of this assumption.

In recent contributions to the directed search literature, the problem posed by the commitment assumption is acknowledged but addressed by apparently only two papers. Masters (2005) builds a directed search model in which advertisements of (only) the legal minimum wage are credible without commitment because firms would incur fines when paying less than the minimum wage. The other is Menzio (2007) where signals fix workers' beliefs prior to wage bargaining so that bargaining leads indeed to higher wages at firms that signalled higher wages, despite the absence of commitment. However, neither discounting nor explicit search costs feature in this model, and the reliance on bargaining makes it largely inapplicable to the directed search literature where firms typically set wages. In conclusion, we would argue that these two papers leave too many questions open to fully close this gap in the literature.

1.3 Thesis overview

This thesis proceeds as follows. Chapter 2 introduces signalling to an environment in which investments in search are not justified despite the existence of gains from trade: a labour market model with frictions where workers search for jobs and perfectly informed firms set wages. For such an environment, Diamond (1971) found equilibrium wages so low that workers cannot even recoup their search costs. Workers in theory then choose not to search at all but do search in practice (the Diamond paradox). We introduce effort and advertisements as choices for workers and firms, respectively. As such, neither resolves the Diamond paradox. The solution we propose combines efficiency wages with wage advertisements. When workers rely on advertisements to decide whether investing in search is justified or not, firms have to advertise wages that allow workers to recoup their search costs. The advertisements fix workers' beliefs. If a firm then reneges on its advertisement the worker's credible retaliation is not to work, i.e. to exert zero effort. Sufficiently patient firms thus find it optimal not to renege but to actually pay the wages that make participation worthwhile for workers. Finally, by requiring contracts with sufficiently high severance payments, workers can ensure that only sufficiently patient firms participate. In equilibrium, workers then do engage in search and may earn positive rents. These results do not depend on a commitment assumption or on any form of competition.

Chapter 3 is relevant for the other question formulated at the beginning of this introduction: it considers a matching market with two-sided heterogeneity where signals help agents to find the best feasible match.[6] Agents can choose a signal of their type and search sequentially. The model incorporates discounting and explicit search costs but does not assume truthful signals. Rather, we identify the condition under which signals are truthful and a unique separating equilibrium with perfect sorting arises despite frictions. We find that supermodularity of the match production function is a necessary and sufficient condition. This is a weaker condition than is needed for sorting in models without signals, which may explain why sorting is much more widespread in reality than existing models would suggest. Supermodularity functions here as both a sorting condition and a single-crossing property. The unique separating equilibrium in our model achieves nearly unconstrained efficiency despite frictions: agents successfully conclude their search after a single meeting, a stable matching results, and overall match output is maximised.

Chapter 3 thus shows that signals in two-sided search can remove the inefficiencies of search models and achieve the optimal outcomes that are normally only achieved

[6]While chapter 3 is written with two-sided matching markets in mind, the model environment features one set of agents who match with each other, which considerably simplifies the analysis.

by matching models. Essentially, agents in our model replace costly search through meetings by costless search through truthful signals, which allows them to avoid almost all search costs. Then the situation effectively approaches a setting without search frictions. That search through signals is effectively costless is therefore central to our argument, and may be thought of as an ideal case that is not yet widely observed in real-world markets. However, as information and communication technologies progress and the costs of search through signals fall, the efficiency gains implied by our model should begin to materialise.

However, a large number of empirical studies suggest that the efficiency of matching has decreased over the last few decades. This widely-held belief is based on the finding of negative time trends in empirical matching functions. In chapter 4, we revisit this puzzling finding and investigate whether it simply arises from omitted variable bias. We consider three kinds of variables that are typically omitted from the matching function: job seekers beyond the unemployed, measures of inflows as opposed to stocks, and agents' selectivity. We first build a model of all labour market flows and use it to construct series for these flows from aggregate data on the U.S. labour market. Using these series, we obtain a measure for employed and non-participating job seekers. When we thus include all job seekers, the estimated time trend remains unchanged. Indirect evidence suggests that the omission of selectivity does not generate the negative time trend either.

We also obtain measures for inflows into unemployment and vacancies. Once these are included, the magnitude of the time trend is halved. The full magnitude of the time trend could only be explained when new results were obtained shortly after the submission of chapter 4, still using the same methods.[7] These results suggest that the failure to account for vacancy dynamics leads to omitted variable bias, which translates into exactly the time trends observed in estimated matching functions (both with or without inflow measures). Hence the evidence from empirical matching functions does not necessarily contradict our theoretical conclusion that signals improve the efficiency of matching.

[7] A version of chapter 4 presenting these results is available as IAB Discussion Paper 3/2012.

Chapter 2

Directed search, efficiency wages and the Diamond paradox

"If your wage is cut a bit, you'd start to misbehave in your job, but you wouldn't leave."

C. A. Pissarides[1]

2.1 Introduction

The most extreme inefficiency that can occur in a market where there are gains from trade is complete market breakdown. This is precisely the fate that Diamond's (1971) results imply for labour markets with frictions where perfectly informed firms set wages and job search is costly for workers. Paradoxically, real-world labour markets often broadly fit this description and yet do not break down. The resolutions of this paradox in the literature appear based on imperfectly informed firms or somehow ensure that firms compete directly for workers despite search frictions. Many other contributions, notably the directed search literature, circumvent the Diamond paradox by ad-hoc assumptions. In this chapter, we propose a solution based on a combination of efficiency wages and advertisements.

[1]On April 9th, 2008 in Birmingham.

The problem. Consider a labour market where firms set wages and job seekers can visit one firm at a time at some search cost. Once the job seeker has come to a firm, her search costs are sunk and do not affect the firm's behaviour. A perfectly informed firm then offers only the worker's reservation wage, that is, a wage as high as unemployment income.[2] Since the job seeker is thus not even reimbursed for the search costs, as a rational forward-looking agent she will choose to avoid these costs by not engaging in search in the first place. As this conclusion extends to any job seeker, there is then no functioning labour market. We refer to this result of theoretical market breakdown as the *Diamond paradox*.[3]

This result is a paradox because many labour markets in the real world do function although Diamond's (1971) setting seems to apply to them: firms set wages and have a good idea about reservation wages, job seekers incur search costs, and yet they engage in search. Hall and Krueger (2008) survey newly hired workers and report that roughly half of these workers were made a wage offer only after their new employer learnt their previous wage. In theory, however, Diamond's (1971) result has proven surprisingly robust, see for example Hopkins and Seymour (2002). Section 3.1 will demonstrate it in very simple terms within our model.[4]

Answers in the literature. An influential article by Burdett and Mortensen (1998) is often mentioned as a solution to the Diamond paradox. However, it is now known that their model of on-the-job search relies heavily on the assumption that firms do not know the employment status of job seekers visiting them; see Postel-Vinay and Robin (2002) and Stevens (2004). Carrillo-Tudela (2009) surveys empirical evidence against this assumption.

Another oft-cited response are efficiency wages, be it from the branch with het-

[2]Diamond's (1971) original model features buyers and sellers, with the finding that the seller can behave as a monopolist and set the profit-maximising price. Provided the seller is perfectly informed, the price is set equal to the buyer's valuation.

[3]In the literature, the degenerate wage distribution implied by Diamond's (1971) results is also often associated with the term "Diamond paradox". We use the definition given by Mortensen and Pissarides (1999): "The Diamond paradox is easily stated: there is no equilibrium in which exchange takes place if the cost of search is strictly positive. Instead, no workers participate." (See Mortensen and Pissarides (1999), p. 2609.) A more elaborate version of this definition is found on p. 2608.

[4]Simple demonstrations in different models can be found in Mortensen and Pissarides (1999) and in Gaumont et al. (2006).

erogeneous workers, be it from the branch that models effort explicitly as a choice variable. From the first branch, particularly Albrecht and Axell (1984) has been associated with responses to the Diamond paradox (although they themselves do not mention the paradox). Yet it is well known now that neither this contribution nor many similar models resolve the paradox.[5] Essentially, Albrecht and Axell (1984) can explain why certain unmatched workers are willing to engage in search *given* that other unmatched workers do so. To solve the Diamond paradox, however, one has to show why these other unmatched workers search in the first place.

From the second branch of efficiency wage models, Shapiro and Stiglitz (1984) are sometimes considered a response to the Diamond paradox. In their model, firms have to pay wages that leave rents to workers because workers would otherwise shirk. However, we show in section 3.4 below that their model can only resolve the Diamond paradox using the assumption of imperfectly informed firms: only if firms are so ignorant of workers' effort that the probability of being caught shirking is sufficiently low, wages to induce effort will be high enough for workers to engage in search. In fact, we find that this kind of efficiency wages as such does not resolve the Diamond paradox either.

Directed search models have introduced wage advertisements into the sort of markets Diamond (1971) considered. Advertised wages allow workers to direct their search to firms advertising wages that at least reimburse them for their search costs. Workers will engage in this directed search - if and only if firms do not renege on their advertisement once a worker has come along. However, we demonstrate in section 3.2 below a clear incentive for firms to renege, so that workers would not engage in any search after all.

To avoid this result, virtually all directed search models assume this problem away by simply saying that firms commit to the wages (or contracts or mechanisms) they advertise.[6] By also leaving unclear how firms might be able to commit, the problem

[5]See Mortensen and Pissarides (1999) and Gaumont et al. (2006).

[6]A contribution by Coles (2001) suggests that firms do not lower the wage ex post because they need to retain their workers. While this is similar in spirit to the ideas developed in this chapter, Diamond's (1971) set-up implies that workers can match at most once and thus cannot gain from quitting a job.

is effectively ignored, with the exception of contributions by Mortensen (2003) and Masters (2005) who explicitly acknowledge the problem. Mortensen (2003) then argues that firms' reputation concerns will lead them not to renege, but there is no room for reputation effects in Diamond's (1971) set-up, nor do real-world workers typically know about a firm's treatment of previous workers. Masters (2005) argues that firms can commit to paying the legal minimum wage if they are otherwise liable to heavy fines. While workers will participate if the minimum wage is sufficiently high, there are nevertheless many labour markets without binding minimum wages, and it remains unclear why workers should participate in these markets. In conclusion, models with just advertisements cannot be said to solve the Diamond paradox.[7]

A number of models avoid the Diamond paradox by allowing for multiple offers to workers as in Burdett and Judd (1983), for example. This way, it is ensured that firms compete at a stage where there are no search frictions for the worker. However, as a worker can only visit one firm at anyone time, this requires that the worker receives or already holds a (potential) job offer from a firm that she is not currently visiting. Then the question arises why firms should not renege on such offers once the worker has come to them. That is, again like advertisements, multiple offers require commitment by firms, without saying how firms can commit.

Ensuring commitment might seem like a purely theoretical problem; after all, workers in the real world do not take glossy advertisements literally. Yet this is easily incorporated into our view: such discrepancies may follow commonly known conventions, so that no-one is misled in equilibrium.[8] The problem arises only when workers encounter unpleasant surprises. A second objection might be that real-world firms rarely make a substantially lower wage offer than was indicated before. While that does not mean that firms have no incentive to do so, it seems that they typically do not follow this incentive. If it was better understood why not, this would point to a solution of the Diamond paradox.

[7]Our criticism of the commitment assumption is in no way qualified by the finding, due to Coles and Eeckhout (2003), that advertising a fixed wage and committing to it constitutes equilibrium behaviour: apart from this being only one of a continuum of equilibria emerging from Coles' and Eeckhout's model, it remains unclear how firms can commit.

[8]Formally, as long as advertisements can be related to firms' actual wage offers by a one-to-one function, the equilibrium will be the same as with truthful advertisements.

Finally, certain labour market institutions trivially resolve the Diamond paradox. For example, unions, binding legal minimum wages, or laws requiring firms to reimburse search costs can ensure that workers engage in search. Apart from violating Diamond's (1971) original assumptions, however, such institutions can only partly explain why labour markets function: many functioning labour markets do not feature any of them.

In conclusion, while both Burdett and Mortensen (1998) and Shapiro and Stiglitz (1984) appear to depend critically on the assumption that firms are imperfectly informed, many other responses to the Diamond paradox seem to depend explicitly or implicitly on the assumption of commitment by firms. Given that neither assumption necessarily holds in real-world labour markets, the question arises whether the Diamond paradox can also be solved when firms are perfectly informed, as they expressly are in Diamond's (1971) original setting, and when they cannot commit.

This chapter. The solution we suggest in this chapter does not violate Diamond's (1971) original assumptions. Instead, we extend the original setting in two notable ways. Firstly, we model the match between firm and worker explicitly as a repeated sequential game in which the firm sets the wage and the worker chooses effort. However, the Diamond paradox is not solved without the second extension, advertised wages (where we exclude competition through advertisements by assumption). This set-up reflects important features of many real labour markets. Hall and Krueger (2008) report from their survey that two thirds of newly hired workers were made a take-it-or-leave-it offer, and that a third had been informed about wages prior to visiting the firm.

Using these extensions, we model how advertisements fix workers' beliefs so that, if the firm reneges, the worker's credible retaliation will be not to work.[9] This may be thought of as "work to rule" or "white strike". Our model thus provides an endogenous explanation why firms do not renege: in equilibrium, sufficiently patient firms advertise wages that make workers' participation worthwhile and optimally

[9]The idea that messages become relevant by fixing beliefs appears to be spreading, but the only other application in the context of search known to us is Menzio (2007).

choose not to renege on them. To ensure that firms are sufficiently patient, workers in our model may require a contract with a severance payment that credibly signals the firm's patience. Hence only sufficiently patient firms will engage in search in equilibrium.

The chapter is organised as follows. Section 2.2 builds the base model. Section 2.3 demonstrates the Diamond paradox within our model and goes on to show that advertisements as such do not resolve it. Section 2.4 proposes a solution by combining advertisements with efficiency wages, under the assumption that firms are sufficiently patient. In this solution, workers do engage in search but do not earn rents. Section 2.5 clarifies that efficiency wages as such would not solve the Diamond paradox. Section 2.6 drops the assumption of patient firms and considers severance payments as a market response to impatient firms. Indeed, section 2.6 finds that this response will in equilibrium play much the same role as the assumption of patient firms, and that workers may even earn rents in equilibrium. Finally, section 2.7 concludes.

2.2 Base model

Our base model is essentially a labour-market version of Diamond's (1971) set-up. Time is discrete with an infinite horizon. The only agents are firms (possibly heterogeneous with respect to time preferences) and homogeneous workers. Each firm is thought of as an organisational unit that can employ at most one worker. Each worker is endowed with an undivisible unit of labour. Firms may match with workers in a marketplace, and matched agents leave the marketplace. There is a possibly large flow of firms with vacancies onto the marketplace and a corresponding flow of unemployed workers. While there are thus flows onto and out of the marketplace, a steady state is not necessary for our analysis.

Following MacLeod and Malcomson (1989), we think of any period t as being divided into subperiods t^0, t^1, and t^2. Consider a period of the matching process on the marketplace that we call period 0. In subperiod 0^0, new firms with vacancies and unemployed workers flow onto the marketplace. Let U_0 denote the number of workers in the marketplace in subperiod 0^1; this number is the sum of those who flowed in and of those who have remained in the marketplace from previous periods. Likewise, V_0 denotes the number of firms in the marketplace in subperiod 0^1. In 0^1, each firm in the marketplace sends one notification of a vacancy to a randomly chosen worker in the marketplace (provided that $U_0 \geq 1$), at explicit search costs $k_F > 0$ to the firm. For the case that a worker receives more than one such notification in subperiod 0^1, we make the following assumption:

Assumption 2.1 (No competition). *Among multiple communications from firms, the worker selects one randomly and ignores all others.*

The expected number $C_0(U_0, V_0)$ of workers contacted through at least one notification follows from an urn-ball process:

$$C_0(U_0, V_0) = U_0 \left(1 - \left(1 - \frac{1}{U_0}\right)^{V_0}\right) \tag{2.1}$$

Then a worker's probability of being contacted in subperiod 0^1 is

$$\eta_0 \equiv \frac{C_0(U_0, V_0)}{U_0} = 1 - \left(1 - \frac{1}{U_0}\right)^{V_0}$$

Workers thus end up with at most one notification and in this case decide whether or not to visit the firm that notified them. Next, in subperiod 0^2, the worker may come to the firm at time-invariant costs $k_W > 0$. Such search costs may be thought of as writing applications or travel costs on the side of workers and recruitment costs on the side of firms. Those not visiting a firm may exit the marketplace and will otherwise begin a period like period 0 anew.

Once the worker has come to a firm at the end of period 0, period 1 will be an interview period for these two agents, and potentially a match. The interview takes place only in subperiod 1^0, as follows. The firm posts a wage w_1 under perfect information:

Assumption 2.2 (Diamond's assumptions). *Firms unilaterally set the wage and observe all relevant variables and parameters.*

If the worker rejects the wage offer, she will return to the marketplace. In this case, 1^1 and 1^2 will be like 0^1 and 0^2, respectively, without any possibility to recall previous offers. If the worker accepts the offer, worker and firm immediately exchange labour for the wage and exit the marketplace. We assume implicitly that the gains from trade are more than enough to reimburse both k_W and k_F. The matched agents do not play a role anymore in 1^1 and 1^2. An interview period can thus turn into either a match or into another period on the marketplace, depending on whether the worker accepts or not.

This structure corresponds to Diamond's (1971) original setting where sellers post prices under perfect information and buyers only learn the actual price after coming to the seller. Buyers then choose whether to buy at this price or continue searching. Diamond's setting is expressly intended for a single purchase of some consumer durable, so that there are no repeat purchases and issues around firms' reputation cannot arise. In our context, this translates into just a single match.

We now specify payoffs. All workers discount the future using the discount factor $\delta_W \in (0,1)$, while firms use discount factor $\delta_F \in (0,1)$. Positive payoffs can be earned only in matches. Let $\Pi_t(w_t)$ be the *value* to the firm of being in a match in period t that pays w_t to the worker. The corresponding value to a worker of being in the match is denoted $\Phi_t(w_t)$. Let $\Omega_{t+1}(w_{t+1})$ be the *expected* value to the worker of learning the actual wage w_{t+1} in an interview at the beginning of next period, having come to a firm. The value to the worker of being in the marketplace in period t is found as

$$\Gamma_t = (1-\eta_t)\delta_W\Gamma_{t+1} + \eta_t \max\left[\delta_W\Omega_{t+1}(w_{t+1}) - k_W, \delta_W\Gamma_{t+1}\right] \tag{2.2}$$

Equation (2.2) says that the worker is not contacted by any firm with probability $1-\eta_t$; when contacted, however, the worker will still not visit the firm if the expected payoff from doing so does not exceed $\delta_W\Gamma_{t+1}$. Whenever the worker does visit a firm and incurs the search costs, she subsequently obtains the expected value of learning the wage,

$$\Omega_{t+1}(w_{t+1}) = \int \max[\Phi_{t+1}(w_{t+1}), \Gamma_{t+1}] d\Lambda(w_{t+1}) \tag{2.3}$$

where $\Lambda(w_{t+1})$ is the known distribution of actual wage offers.[10] That is, she will reject the wage offer if she prefers to continue the period unmatched rather than matched at a wage w_{t+1}.

2.3 The Diamond paradox with advertisements

2.3.1 A simple demonstration of the Diamond paradox

Suppose a worker has come to a firm at the end of period t and is in an interview in period $t+1$. By rejecting the wage offer, she would obtain Γ_{t+1}. The firm seeks to minimise its wage costs by offering a wage that leaves the worker indifferent between accepting and rejecting (or an ϵ more to tip the balance), so that

$$\Phi_{t+1}(w_{t+1}) = \Gamma_{t+1}$$

in equation (2.3), implying

$$\Omega_{t+1}(w_{t+1}) = \Gamma_{t+1}$$

Then equation (2.2) becomes

$$\Gamma_t = (1-\eta_t)\delta_W \Gamma_{t+1} + \eta_t \max\left[\delta_W \Gamma_{t+1} - k_W, \delta_W \Gamma_{t+1}\right]$$

which simplifies to

$$\Gamma_t = \delta_W \Gamma_{t+1}$$

This means that the worker simply chooses not to visit the firm - after all, the only difference a visit makes is that the worker incurs the search costs k_W. Since exactly the same logic obtains in the next period, the worker again decides against the visit, hence $\Gamma_{t+1} = \delta_W \Gamma_{t+2}$, and so on. The worker always continues to search, never

[10] In all equilibria that we consider, $\Lambda(w_{t+1})$ consists of a single wage. Forward-looking workers can therefore easily infer it from the behaviour of individual firms.

matches, and never earns any payoff, so that $\Gamma_t = \Gamma_{t+1} = \Gamma_{t+2} = \ldots = 0$. This argument generalises to any worker and any interview period. Note that it does not depend on the value of η_t, so that even a single worker without competition from others (i.e. $\eta_t = 1$) would obtain a zero payoff.

In conclusion, no worker actually engages in search, as the search costs would never be recouped. This result of an inactive market as the unique equilibrium when there are search costs (or entry costs) is the Diamond paradox. The root cause of the Diamond paradox are thus search costs that are incurred ex ante and become sunk costs by the time a worker visits a firm, so that the firm has no reason to reimburse the worker for these costs.

2.3.2 Introducing advertisements

If search costs that are sunk lie at the heart of the Diamond paradox, wage advertisements appear promising as a remedy: if firms advertise wages to workers before they visit a firm, this may bridge the step that makes these costs sunk costs. Workers would then only visit firms whose advertised wage offer suggests that they are reimbursed for their costs. Then only such wage offers attract applicants, which forces firms to advertise such wage offers. This argument, however, critically depends on advertisements being truthful, so that no firm reneges on its advertised wage offer and actually makes a different offer once a worker has come along.

We can formally evaluate this argument by replacing the mere notifications about vacancies in our set-up by advertisements that include a precise wage figure, denoted $\tilde{w}$. To focus on advertised wages as such, we apply assumption 2.1 also here, so that workers select an advertisement randomly and firms therefore cannot compete in advertisements. Suppose that workers take $\tilde{w}$ at face value. Then we have *directed search* at least in the sense that workers can direct their search away from firms whose $\tilde{w}$ would not reimburse workers for search costs. Denote by $\Omega_{t+1}(w_{t+1}|\tilde{w}_{t+1})$ the expected value to the worker of learning the offered wage in an interview, given that $\tilde{w}_{t+1}$ has been advertised. Since workers take advertisements at face value,

$$\Omega_{t+1}(w_{t+1}|\tilde{w}_{t+1}) = \max[\Phi_{t+1}(\tilde{w}_{t+1}), \Gamma_{t+1}]$$

The best outcome for the firm is that the worker comes to the firm and then works at the lowest acceptable wage. At time t, the firm's optimal choice of $\tilde{w}_{t+1}$ therefore ensures

$$\begin{aligned} \delta_W \Omega_{t+1}(w_{t+1}|\tilde{w}_{t+1}) - k_W &\geq \delta_W \Gamma_{t+1} \\ \Leftrightarrow \quad \delta_W \max[\Phi_{t+1}(\tilde{w}_{t+1}), \Gamma_{t+1}] - k_W &\geq \delta_W \Gamma_{t+1} \end{aligned} \tag{2.4}$$

Equation (2.4) can only possibly be satisfied if $\Phi_{t+1}(\tilde{w}_{t+1}) > \Gamma_{t+1}$. By contrast, the firm's optimal choice of w_{t+1} sets, as before,

$$\Phi_{t+1}(w_{t+1}) = \Gamma_{t+1}$$

which implies

$$\Phi_{t+1}(\tilde{w}_{t+1}) > \Phi_{t+1}(w_{t+1}) \quad \Rightarrow \quad \tilde{w}_{t+1} > w_{t+1}$$

In words, the firm's optimal behaviour is *time inconsistent*: it would like to induce worker participation through advertisements, but once workers participate it reneges on these advertisements. Rational workers would be aware of this inconsistency, would therefore distrust firms' advertisements, and would not participate. The result is still the Diamond paradox. Therefore, the numerous contributions that simply assume firms commit to their advertisements essentially assume the Diamond paradox away.

2.4 A solution: advertised efficiency wages

2.4.1 Extended model

In this section, we also introduce effort as a choice variable for workers in addition to firms' advertisements. In order to elicit effort from workers during the match, firms might pay efficiency wages,[11] and we will show that the combination of efficiency wages and advertisements resolves the Diamond paradox. Section 2.5 will then clarify that it is the combination that delivers the result, since efficiency wages as such do not solve the Diamond paradox.

We extend the base model with advertisements by viewing the match as repeated and dynamic productive interaction between worker and firm, rather than a one-off exchange transaction. Production technology is identical across firms and uses labour as the only input. The worker produces perfectly measurable, yet not verifiable, output $p(e_t)$ per period during the match, where the continuous $e_t \geq 0$ denotes effort in period t. Let the price for output be 1, so that $p(e_t)$ is also the monetary value of output. From exerting effort, a worker incurs costs $c(e_t)$ measured as the monetary value of her disutility from effort. We assume that these costs are convex, that output per worker is concave, and that there are gains from trade:

Assumption 2.3 (Regularity conditions). *Output per worker $p(\cdot)$ satisfies $p'(e) > 0$, $p''(e) < 0$, and $p(0) = 0$, while the costs of effort $c(\cdot)$ satisfy $c'(e) > 0$, $c''(e) > 0$, and $c(0) = 0$. Finally, $\exists\ e > 0$ such that*

$$\frac{\min[\delta_W, \delta_F][p(e) - c(e)]}{1 - \min[\delta_W, \delta_F]} > k_W + k_F \tag{2.5}$$

Equation (2.5) says that both the firm's and the worker's search costs can be recouped by the present discounted value of match production over the infinite horizon, where we discount to the period 0 just before the match because search costs are incurred in that period.

[11] Models with effort constitute one major branch of efficiency wage models. The other branch, as in Weiss (1980), features differential reservation wages, which we mentioned in the context of Albrecht and Axell (1984).

The timing in period 0 stays as in the base model, but firms now advertise wage information instead of merely informing workers of a vacancy. As all communications by firms are now advertisements, by assumption 2.1 workers select an advertisement at random. How the worker responds to the selected advertisement will depend on her beliefs. The interview period changes as follows. In 1^0, the firm does not only make a wage *offer*, denoted $\hat{w}_1 \geq 0$, but also demands a level of e_1, which is not modelled explicitly. The worker accepts or rejects. If she rejects, she will return to the marketplace as before, so that 1^1 and 1^2 will be like 0^1 and 0^2. If she accepts, the match is concluded and neither worker nor firm can return to the marketplace in 1^1. Instead, the worker then freely chooses effort e_1 and produces for the firm in 1^1. The firm pays the actual wage w_1 only in 1^2.

Any period $t \geq 2$ is then a period of a continued match. The actions in this period and their timing are exactly as in the interview period, with the exception that a rejection by the worker now does not end an interview, but terminates the entire match. (The firm can thus terminate the match by provoking a rejection by the worker, through its current wage offer or its behaviour in a previous period.) To summarise:

Period 0: on the marketplace

0^0: new workers and firms flow in

0^1: firms contact workers through advertisements and incur k_F; the contacted decide whether to visit a firm

0^2: workers who so decided come to a firm and incur k_W

Period 1: interview and potentially first period of a match

1^0: firm offers wage $\hat{w}_1$; worker accepts or rejects

1^1: workers who accepted work; others return to search and pass a subperiod like 0^1

1^2: workers who accepted receive w_1; others pass a subperiod like 0^2

The match lasts forever unless either side chooses to terminate it. While agents may terminate the match, they cannot return to a previous stage. As before, only a single match is thus possible for each agent. Note that each period of a match constitutes an extensive game where agents move sequentially in subperiods. The match is thus a potentially infinite repetition of this extensive stage game. In each match period t, their actions $\hat{w}_t \in \mathbb{R}_+$, $e_t \in \mathbb{R}_+$, and $w_t \in \mathbb{R}_+$ combine to an outcome $(\hat{w}_t, e_t, w_t) \in \mathbb{R}^3_+$. The firm's preference relation on $\mathbb{R}^3_+$ is represented by its payoff function $\pi : \mathbb{R}^3_+ \to \mathbb{R}$ (i.e. its profit function) so that a weak preference for outcome $(\hat{w}_t, e_t, w_t)$ over $(\hat{w}'_t, e'_t, w'_t)$ implies

$$\pi(\hat{w}_t, e_t, w_t) \geq \pi(\hat{w}'_t, e'_t, w'_t)$$

The worker's preference relation on $\mathbb{R}^3_+$ is similarly represented by a payoff function $b : \mathbb{R}^3_+ \to \mathbb{R}$. We specify these payoff functions in a very natural way as

$$\pi(\hat{w}_t, e_t, w_t) = p(e_t) - w_t \quad \text{and} \quad b(\hat{w}_t, e_t, w_t) = w_t - c(e_t)$$

so that $b_t \equiv b(\hat{w}_t, e_t, w_t)$ denotes the part of the wage over and above $c(e_t)$. The firm's preferences in the infinitely repeated game are defined on the set of infinite sequences of stage game outcomes and satisfy weak separability. Concretely, they are given by δ-discounting, i.e. an infinite sequence of payoffs $(\pi(\hat{w}_t, e_t, w_t))_{t=1}^{\infty}$ is weakly preferred to an infinite sequence $(\pi(\hat{w}'_t, e'_t, w'_t))_{t=1}^{\infty}$ if

$$\sum_{t=1}^{\infty} \delta_F^{t-1} [\pi(\hat{w}_t, e_t, w_t) - \pi(\hat{w}'_t, e'_t, w'_t)] \geq 0$$

and equivalently for the worker's preferences.

We would argue that firms in the real world advertise both the total wage they offer and the effort they expect. Of course, such advertisements imply an advertised b_t as the advertised wage less the disutility of the expected effort. Since workers in our model only care about the sequence of b_t, we can pretend for simplicity that firms' only advertise information on b_t. Strictly speaking, a firm's advertisement concerns the infinite sequence $(b_t)_{t=1}^{\infty}$, but it will turn out that firms choose to advertise a time-invariant level of b_t to be paid in each period forever. Hence we simply let $\tilde{b}_t$

denote this level of b_t without loss of generality.

The histories in subperiods t^0, t^1, and t^2 can then be expressed as

$$\begin{aligned} h(t^0) &= \{\tilde{b}, (\hat{w}_1, e_1, w_1), (\hat{w}_2, e_2, w_2), \ldots, (\hat{w}_{t-1}, e_{t-1}, w_{t-1})\} \\ h(t^1) &= h(t^0) \cup \{\hat{w}_t\} \\ h(t^2) &= h(t^1) \cup \{e_t\} \end{aligned}$$

with $h(1^0) = \{\tilde{b}\}$. Finally, let us recursively define $\Pi_t(e_t, w_t)$ as the value to the firm of being in a match in period t:

$$\Pi_t(e_t, w_t) = \max[p(e_t) - w_t + (1 - \alpha_t)\delta_F \Pi_{t+1}(e_{t+1}, w_{t+1}), 0] \tag{2.6}$$

where α_t is the probability that the worker terminates the match before period $t+1$. Equation (2.6) says that a firm in a match can choose to continue the match, earning $p(e_t) - w_t$ in t.[12] Then the firm will enter the same situation in period $t+1$ provided the worker does not terminate the match, and will otherwise be unmatched at the beginning of period $t+1$, earning 0 forever. Alternatively, the firm can itself terminate the match and obtain 0. Similarly, let $\Phi_t(e_t, w_t)$ be the value to the worker of being in a match in period t:

$$\Phi_t(e_t, w_t) = \max[w_t - c(e_t) + (1 - \beta_t)\delta_W \Phi_{t+1}(e_{t+1}, w_{t+1}), 0] \tag{2.7}$$

where β_t is the probability that the firm terminates the match before period $t+1$.

2.4.2 Equilibrium: endogenous commitment

We first focus on the stage game in some period t of a match between a firm and a worker. Analysing this game in isolation means treating it like a one-shot game without any repetition. Then the relevant payoffs are just the payoffs from period t (not counting for the moment any sunk costs). In analysing the stage game, which is

[12] The formulation of equation (2.6) encompasses the case that the worker terminates the match already in period t: then $e_t = p(e_t) = 0$ and $\alpha_t = 1$.

an extensive game, we employ the concept of a subgame perfect equilibrium (SPE). We find that the unique SPE of the stage game is degenerate, a now well-known result (see Malcomson (1999)) that is proven by backward induction.

Lemma 2.1 (One-shot game). *In the unique SPE of the one-shot game,* $e_t^* = w_t^* = 0$ *(and* $\hat{w}_t^*$ *indeterminate).*

Proof. In period t^2, the firm takes equilibrium actions $\hat{w}_t^*$ and e_t^* as given because they have been taken in periods t^0 and t^1. As $\pi(\hat{w}_t, e_t, w_t|\hat{w}_t, e_t)$ is then unambiguously falling in w_t, the firm optimally sets $w_t^* = 0$. By backward induction, the worker in subperiod t^1 knows she will receive a wage 0 in any case, and thus maximises $b(\hat{w}_t, e_t, w_t|\hat{w}_t, w_t = 0)$ by choosing $e_t^* = 0$. All this is independent of the firm's choice of $\hat{w}_t$ in t^0, so that $\hat{w}_t$ may take any value in $\mathbb{R}_+$ in this equilibrium. □

We are thus left with a situation of no production: $\pi(\hat{w}_t^*, e_t^*, w_t^*) = b(\hat{w}_t^*, e_t^*, w_t^*) = 0$. Turning to the SPE of the infinitely repeated game in matches, the following SPE results when players in a match maximise their payoffs period by period:

Lemma 2.2 (Worst SPE). $\exists$ *a SPE of the infinitely repeated game such that* $e_t^* = w_t^* = 0$ *(and* $\hat{w}_t^*$ *indeterminate) in every period* t *forever and hence* $\Phi_t(e_t^*, w_t^*) = \Pi_t(e_t^*, w_t^*) = 0$*.*

Proof. By definition, a SPE must also be a NE because any game is a proper subgame of itself. The SPE of the stage game identified in lemma 2.1 is therefore also a NE, and the infinite repetition of a NE of the stage game must be a SPE of the infinitely repeated game. It remains to verify that the stage game is in fact repeated infinitely under these circumstances. As $e_t = 0$ in every period t while the match lasts, no surplus is ever generated so that $\Phi_t(e_t^*, w_t^*) = \Pi_t(e_t^*, w_t^*) = 0$, yet terminating the match also only carries a payoff of 0 for each player. Hence neither player prefers to terminate the match, $\alpha_t^* = \beta_t^* = 0 \; \forall t$, and the stage game is repeated infinitely. □

Hence a situation of no production each period constitutes a SPE. To prepare later results, we note that neither lemma 2.1 nor lemma 2.2 depend at all on the value of the firm's discount factor δ_F. Of course, this equilibrium is unsatisfactory for the players, given that, by assumption 2.3, a different feasible choice of e_t would lead to a positive surplus that could then be shared in a Pareto-improving way. The firm can achieve this as follows. Consider a match where the firm pays w_t only in return for a certain choice of e_t. The worker will not choose e_t unless she expects the wage to at least cover $c(e_t)$, hence $w_t \geq c(e_t)$ would be necessary, i.e. $b_t \geq 0$. Because the firm could always pay some ϵ greater than but very close to 0 in order to tip the balance, we henceforth suppose that the worker chooses e_t as demanded by the firm in case of indifference. Since the firm can freely choose which e_t to reward, it will each period choose the e_t that maximises its payoff:

$$e_t^* \equiv \arg\max_{e_t} \pi(\hat{w}_t, e_t, w_t) = \arg\max_{e_t} p(e_t) - c(e_t) - b_t \tag{2.8}$$

will hold in such an equilibrium. Whenever b_t does not depend on e_t, the firm will choose a (time-invariant) level we denote by $\bar{e}$:

$$\bar{e} \equiv \arg\max_{e} p(e) - c(e) \tag{2.9}$$

The next lemma claims that $e_t = \bar{e}$ can indeed arise as an equilibrium action in each period of the infinitely repeated game, provided the firm is patient enough.

Lemma 2.3 (Cooperative SPE). *For every $b_t \in [0, p(\bar{e}) - c(\bar{e}))$, $\exists\ \delta_F^* < 1$ such that for all $\delta_F \in [\delta_F^*, 1)$, the profile $(\hat{w}_t, e_t, w_t) = (c(\bar{e}) + b_t, \bar{e}, c(\bar{e}) + b_t)$ in every match period t is a SPE of the infinitely repeated game. The worker and the firm respectively obtain payoffs $\Pi_t(e_t^*, w_t^*) > 0$ and $\Phi_t(e_t^*, w_t^*) \geq 0$.*

Proof. Choose and fix some $b_t \in [0, p(\bar{e}) - c(\bar{e}))$ to be paid in every period. Let players use *grim trigger strategies* as follows. In each period t^0, the firm offers $\hat{w}_t = c(\bar{e}) + b_t$ and demands $e_t^* = \bar{e}$ if there has been no earlier deviation by either side. In each period t^1, the worker chooses $e_t = \bar{e}$ unless there was a deviation. In each period t^2 the firm pays $c(\bar{e})+b_t$ unless there was a deviation. If there was a deviation by either side, the worker instead chooses $e_t = 0$, the firm chooses $w_t = 0$ (with $\hat{w}_t$ indeterminate), and likewise in every subsequent period. That is, players switch to the SPE in lemma 2.2 in case of a deviation.[13] As noted before, any SPE must also be a NE, and therefore the threat of switching is credible.

These threats can thus sustain a SPE with $e_t^* = \bar{e}$ and $w_t = c(\bar{e})+b_t$ if both players prefer not to deviate, given that the threats would be carried out. The worker will obtain $b_t/(1-\delta_W) \geq 0$ if there is never a deviation, but would obtain 0 after a deviation, and hence does not have an incentive to deviate. As to the firm, when there is no deviation it obtains

$$\sum_{t=0}^{\infty} \delta_F^t [p(\bar{e}) - c(\bar{e}) - b_t] > 0 \quad \forall b_t \in [0, p(\bar{e}) - c(\bar{e}))$$

where the inequality follows because assumption 2.3 implies

$$p(\bar{e}) - c(\bar{e}) > 0$$

Hence $\Pi_t(e_t^*, w_t^*) > 0$ here. If instead the firm deviates (i.e. it pays nothing at all), it obtains a short-term gain of $c(\bar{e})+b_t$ in the period of deviation and then obtains 0 in the SPE in lemma 2.2. The short-term gain will be outweighed by the loss sustained from switching to the SPE in lemma 2.2 if

$$c(\bar{e}) + b_t \leq \delta_F \Pi_{t+1}(e_{t+1}^*, w_{t+1}^*) = \sum_{t=1}^{\infty} \delta_F^t [p(\bar{e}) - c(\bar{e}) - b_t]$$

where, following the one-deviation principle, we do not consider a further deviation. This condition holds with equality at δ_F^*:

$$c(\bar{e}) + b_t = \frac{\delta_F^* [p(\bar{e}) - c(\bar{e}) - b_t]}{1 - \delta_F^*} \quad \Rightarrow \quad \delta_F^* = \frac{c(\bar{e}) + b_t}{p(\bar{e})} \tag{2.10}$$

[13]The only rational direction of a deviation here is downward. If, however, the firm was for some reason nice enough to suddenly pay more, this would also be treated as a deviation. Such anomalies arise frequently in this kind of SPE.

That is, δ_F^* is the lowest discount factor the firm may have for this SPE to be sustainable. Since $b_t \geq 0$, $\delta_F^* > 0$. Since $b_t < p(\bar{e}) - c(\bar{e})$, $\delta_F^* < 1$. Hence both players do not deviate if the firm is sufficiently patient. Finally, $\Pi_t(e_t^*, w_t^*) \geq 0$ and $\Phi_t(e_t^*, w_t^*) \geq 0$ imply that players also do not terminate the match, so that $\alpha_t^* = \beta_t^* = 0$, $\forall t$. $\square$

To illustrate, suppose for a moment that firms' δ_F just satisfies equation (2.10) if, say,

$$b_t = \frac{1}{2}\left[p(\bar{e}) - c(\bar{e})\right]$$

There is in fact a continuum of SPE in which firms pay some b_t between (and including) 0 and $(1/2)\left[p(\bar{e}) - c(\bar{e})\right]$ while workers choose $e_t = \bar{e}$. In any such SPE with $b_t > 0$, firms pay more than the minimum needed to induce workers to choose $\bar{e}$, but if a firm unilaterally reduces b_t, this will be a deviation from equilibrium play. The worker would respond to the deviation by henceforth choosing $e_t = 0$, which is credible because this action is in line with a NE. This threat sustains a $b_t > 0$ in SPE provided δ_F is high enough. Even a b_t very close to $p(\bar{e}) - c(\bar{e})$ may be sustained if δ_F is very close to 1, but $b_t = p(\bar{e}) - c(\bar{e})$ would mean that the gain to the firm from not deviating is 0 while there is always a short-term gain from deviating.

It is not crucial for the proof of lemma 2.3 that play moves to the worst SPE (see lemma 2.2) after a deviation. It is crucial, however, that play moves to some SPE in which the worker works sufficiently less so that this SPE is worse for the firm. Any such SPE with lower e_t will be equivalent, after an appropriate normalisation, to the worst SPE. In fact, we think of the worst SPE as "work to rule" or "white strike" instead of literally zero work.[14]

Given the results from lemmas 2.1 to 2.3 for interaction in matches, we now turn to the equilibrium that includes advertising and search behaviour. Consider a worker who is contacted through advertisements (and randomly selects one among them).

[14] To skeptics of game theory, the move all the way to the worst SPE might seem mechanical: can there not be a "compromise" SPE in which the firm pays some lower but still positive b_t and the worker works? While such an objection might be made to a great many SPE in the literature, the response appears particularly simple in this chapter. Given that it deviates, a profit-maximising firm would want to pay the smallest b_t that is still positive, which does not exist: for any small but positive b_t, one can find a lower b_t that is also still positive. Hence a "compromise" SPE does not exist.

Let $\mu(b_t = \tilde{b}|\tilde{b})$ denote the worker's subjective probability that the SPE with $b_t = \tilde{b}$ will be played during the match when $\tilde{b}$ has been advertised. We call an advertisement $\tilde{b}$ *potentially truthful* if there exists a SPE in which the firm pays $b_t = \tilde{b}$. Let us further define $\underline{b}$ and $\bar{b}$ such that

$$\frac{\delta_W \underline{b}}{(1-\delta_W)} = k_W \quad , \quad \frac{\delta_F[p(\bar{e}) - c(\bar{e}) - \bar{b}]}{1-\delta_F} = k_F \tag{2.11}$$

so that every $b_t \in [\underline{b}, \bar{b}]$, if paid in each period forever, would allow both the worker and the firm to recoup their respective search costs.[15] By assumption 2.3, $\underline{b} < \bar{b}$ and each firm is thus prepared to pay $b \in [\underline{b}, \bar{b}]$ if necessary. For reasons that will become clear, we also assume for now that firms are sufficiently patient to pay at most some $B \in [\underline{b}, \bar{b}]$ during the match:

Assumption 2.4 (Patient firms). *Firms' common δ_F just satisfies equation (2.10) when $b_t = B$, for some $B \in [\underline{b}, \bar{b}]$.*

Then the model has a *perfect Bayesian equilibrium* (PBE) with all firms paying $b_t = \underline{b}$ in each match period, as follows.

Proposition 2.1 (PBE with worker participation). *$\exists$ a PBE of the model with advertisements in which*

(i) each firm chooses the advertisement $\tilde{b} = \underline{b}$

(ii) workers trust all potentially truthful advertisements: $\mu(b_t = \tilde{b}|0 \leq \tilde{b} \leq B) = 1$ and $\mu(b_t = \tilde{b}|\tilde{b} > B) = 0$

(iii) the SPE profile $(\hat{w}_t, e_t, w_t) = (c(\bar{e}) + \underline{b}, \bar{e}, c(\bar{e}) + \underline{b})$ is played, $\forall t$ of a match

(iv) workers engage in search and match at the first opportunity.

Proof. Given workers' beliefs as specified under *(ii)*, workers who select an advertisement $\tilde{b} < \underline{b}$ would trust it and consequently would not visit this firm. Advertisements $\tilde{b} > B$ are not trusted. Thus only advertisements $\tilde{b} \in [\underline{b}, B]$ attract workers.

[15] As before, we discount earnings in a match period to the period preceding the match, since agents incur k_W and k_F in this period.

Given their beliefs, workers visiting a firm treat it as a deviation from the SPE they expected when a firm reneges on its advertisement. Let players use grim trigger strategies analogous to those in the proof of lemma 2.3. If a firm reneges, play therefore moves to the SPE in lemma 2.2, giving both players payoff 0. Firms will thus be strictly better off if a SPE as in lemma 2.3 is played with $b_t \in [\underline{b}, B]$ because, by equation (2.11), the firm obtains a payoff $k_F > 0$ even with $B = \bar{b}$. By assumption 2.4, the firm is also patient enough to sustain a SPE with $b_t \in [\underline{b}, B]$. Hence, given workers' reactions to reneging, firms strictly prefer not to renege.

Among all combinations of some $b_t \in [\underline{b}, B]$ and its truthful advertisement $\tilde{b}$ that still attracts workers, profit-maximising firms will prefer the one with the lowest b_t, which they can choose by advertising accordingly. Hence the SPE with $b_t = \underline{b}$ is always played and preceded by the truthful advertisement $\tilde{b} = \underline{b}$. Since $\underline{b}$ does not depend on e_t, the firm demands $e_t^* = \bar{e}$. As all firms advertise and pay $\underline{b}$, there is nothing to gain from visiting a second firm, and the worker thus matches at the first opportunity and incurs k_W only once. By equation (2.11), the payoffs from the SPE with $b_t = \underline{b}$ allow workers to exactly recoup k_W. In terms of equation (2.2), we thus have

$$\delta_W \Omega_{t+1}(w_{t+1} | \tilde{b} = \underline{b}) = k_W$$

Now suppose

$$\delta_W \Gamma_{t+1} > \delta_W \Omega_{t+1}(w_{t+1} | \tilde{b} = \underline{b}) - k_W = 0 \tag{2.12}$$

in equation (2.2). Then $\Gamma_t = \delta_W \Gamma_{t+1}$ and likewise in all subsequent periods. Yet if the worker thus never visits a firm and never earns any payoffs, equation (2.12) cannot hold. Hence

$$\delta_W \Omega_{t+1}(w_{t+1} | \tilde{b} = \underline{b}) - k_W \geq \delta_W \Gamma_{t+1}$$

so that workers engage in search. Since $\underline{b} < \bar{b}$, firms also recoup their search cost. As workers are homogeneous, firms also match at the first opportunity.

Finally, analogous grim trigger strategies would sustain every other SPE with $b_t \in [0, B]$, so that workers correctly believe that they can trust all potentially truthful advertisements (while they would only visit firms advertising at least $\underline{b}$). As firms are not sufficiently patient to pay some $b_t > B$, workers also correctly believe that such SPE will be played with 0 probability. $\square$

2.4.3 Discussion

The cheap-talk PBE in proposition 2.1 can be viewed as one of the SPE in lemma 2.3 extended by advertisements and by the beliefs such advertisements can induce. The key result is that workers engage in search. Advertisements are crucial here because they serve as a benchmark against which a deviation by the firm is defined. Once advertisements have fixed workers' beliefs, the same credible "punishment" for a deviating firm in a SPE in lemma 2.3 now also applies to firms reneging on their advertisements. Intuitively, one could say that the firm loses the worker's goodwill if it reneges on its advertisement. Then the firm prefers not to renege and actually pays the advertised wage. This wage must be sufficient to reimburse the worker for her search costs because she would otherwise not have visited the firm with this advertisement.[16]

As an illustration, consider what would happen if a firm tried to play the SPE in lemma 2.3 with $b_t = 0$. If the firm advertises $\tilde{b} = 0$ (or any other $\tilde{b} < \underline{b}$), no worker will visit. The same holds for advertisements greater than B. If the firm advertises some $b_t \in [\underline{b}, B]$, workers will expect a SPE with this b_t. Then any attempt by the firm to play the SPE with $b_t = 0$ is seen by the worker as a deviation. Concretely, when the worker has to choose e_1, she can compare the advertised $\tilde{b}$ to the wage offer $\hat{w}_1$, having also been told that effort level $\bar{e}$ is expected. If $\hat{w}_1 < c(\bar{e}) + \tilde{b}$, the firm reneges on the advertisement and the worker treats this as a deviation from the SPE in which $\hat{w}_1 = c(\bar{e}) + \tilde{b}$ would have been equilibrium play.[17] Consequently, the worker chooses $e_t = 0$ in every period t of the match, leading to the SPE in lemma 2.2. Hence

[16] Any babbling equilibrium would not fix workers' beliefs about the SPE that is played in the match. Instead, workers might believe that the SPE with $b_t = 0$ arises with a certain probability, regardless of the advertisement. If firms can thus renege without consequences, all firms will do so. As workers' beliefs must be correct in a PBE, the only beliefs workers may have in the babbling equilibrium are $\mu(b_t = \tilde{b} | \tilde{b} > 0) = 0$. Hence workers prefer not to participate, so that the babbling equilibrium leads to the Diamond paradox, in contrast to the PBE in proposition 2.1.

[17] $\hat{w}_1 = c(\bar{e}) + \tilde{b}$ is equilibrium play in the SPE with $b_t = \tilde{b}$ but also in the worst SPE where $\hat{w}$ is indeterminate (see lemma 2.2). However, the advertisement leads the worker to believe that the former SPE is being played when choosing e_1, which is in this context advantageous to the firm.

the firm can only choose between outcomes with payoff 0 (when no worker visits or when the SPE in lemma 2.2 results) and some SPE with $b_t \in [\underline{b}, B]$.

In conclusion, this section has formalised the idea that workers who are confronted with a unilateral reduction of the promised wage can react by "work to rule", working as little as possible. This reaction is consistent with equilibrium play in a subgame-perfect equilibrium and the threat of this reaction therefore induces firms not to renege on their advertisements. Truthful advertisements in turn will only attract workers if the wages reimburse workers for their search costs, so that workers do engage in search. Since workers by assumption select one firm's notification randomly, competition does not play a role in obtaining this result. Imperfect information does not play any role here either, nor does wage dispersion: with $w_t = c(\bar{e}) + \underline{b}$ in every period t, the equilibrium wage distribution consists of a single wage. Wage dispersion would be generated in our model if we relaxed assumption 2.1 to allow for a minimum of competition, for example as in Halko et al. (2008). What does play a role for the results in this section is that firms are sufficiently patient by assumption 2.4, and section 2.6 will address this issue.

2.5 The role of efficiency wages

This section takes a closer look at the role that efficiency wages play for the solution to the Diamond paradox we proposed above. Let us consider the extended model without advertisements, so that firms merely notify workers of vacancies:

$$h(t^0) = \{(\hat{w}_1, e_1, w_1), (\hat{w}_2, e_2, w_2), \ldots, (\hat{w}_{t-1}, e_{t-1}, w_{t-1})\}$$

with $h(1^0) = \emptyset$. Without an advertisement that fixes the worker's beliefs, it is not clear at all which SPE might be played in a match with a firm (that is sufficiently patient not to deviate). In particular, the worker has no reason to expect that SPE with $b_t > \underline{b}$ will be played: lemma 2.3 has established that $b_t \in [0, \underline{b})$ can just as well arise together with $e_t = \bar{e}$ as equilibrium actions in every period. Even when $b_t = 0$ the worker does not have an incentive to deviate, as she obtains $\Phi_t(e_t^*, w_t^*) = 0$

whether or not she deviates.

Therefore, workers have no reason to expect SPE payoffs in this set-up that allow them to recoup their search costs. From the extended model without advertisements, we thus could not conclude that workers engage in search.[18] Section 2.3 has found that advertisements as such also do not ensure that workers engage in search. Hence it is the combination of efficiency wages and advertisements that drives the solution in section 2.4, while *efficiency wages as such do not solve the Diamond paradox.*

How can this conclusion be reconciled with the wide-spread view that efficiency wages in Shapiro and Stiglitz (1984) solve the Diamond paradox? In their model, firms pay indeed higher wages, equivalent to $b_t > 0$ here. Yet their result is driven by their assumption of imperfect information for the firm: it does not always observe the worker's effort. To see this, suppose a worker who "shirks" by choosing $e_t = 0$ is only caught with probability $q < 1$ (thus far, our model has implicitly assumed $q = 1$). Then the SPE in lemma 2.3 with $b_t = 0$ is no longer feasible: by choosing $e_t = \bar{e}$ as before, a worker would obtain payoff 0, but by choosing $e_t = 0$ in anyone period she will in expectation obtain a strictly positive payoff $(1-q)c(\bar{e})$. That is, she will be laid off if caught but will otherwise be paid a wage $w_t = c(\bar{e})$ that compensates her for effort she has never made. To ensure that the worker chooses $e_t = \bar{e}$ under these circumstances, the payoff from working needs to outweigh the expected payoff from shirking. Thus the no-shirking condition is

$$b_t + \delta_W \Phi_{t+1}(e_{t+1}, w_{t+1}) \geq (1-q)[c(\bar{e}) + b_t + \delta_W \Phi_{t+1}(e_{t+1}, w_{t+1})]$$

To limit wage costs as far as possible while still incentivising workers, firms will pay the lowest b_t that satisfies this condition, denoted b^{NSC}. It is determined as the b_t at which the condition holds with equality, so that the worker's payoff from working would be the same as from shirking. Therefore, $\Phi_{t+1}(e_{t+1}, w_{t+1})$ is found as the value of earning b^{NSC} forever, $b^{NSC}/(1-\delta_W)$. With this, we can solve the no-shirking

[18] Section 2.4 implicitly assumes that no firm strictly prefers to send mere notifications of vacancies to workers instead of advertisements. This does not necessarily involve a loss of generality, as mere notifications might fail to attract any workers, while workers were shown to be attracted by sufficiently promising advertisements.

condition for b^{NSC}:

$$b^{NSC} = (1-\delta_W)\frac{1-q}{q}c(\bar{e})$$

Firms will thus pay a wage that reimburses workers for their search costs if and only if the value of earning b^{NSC} over the entire match duration weakly exceeds k_W:[19]

$$\frac{\delta_W b^{NSC}}{1-\delta_W} \geq k_W \quad \Leftrightarrow \quad \delta_W \frac{1-q}{q}c(\bar{e}) \geq k_W$$

This will hold for q low enough but cannot hold for $q = 1$. Hence the assumption of imperfectly informed firms drives the finding of higher wages in Shapiro and Stiglitz (1984). While the imperfections they assume appear much more plausible to us than the information imperfections assumed by Burdett and Mortensen (1998), it remains an assumption that violates Diamond's (1971) original set-up and that does not always hold in reality, as effort can be observed well in many industries (e.g. construction). Instead we have proposed a solution to the Diamond paradox that adds to or refines Diamond's (1971) settings but does not violate them.

2.6 Impatient firms and severance payments

This section explores what happens if assumption 2.4 fails and also firms participate that are not sufficiently patient. We show that contracts with severance payments as an institution then allow patient firms to distinguish themselves from impatient firms. In this context, we find that wages may rise even above the level that exactly compensates workers for their search costs. However, a second finding is that firms will prefer inefficiently low levels of effort in order to limit this rise in wages. To begin, let us drop assumption 2.4 and adopt instead the following assumption:

Assumption 2.4' (Patient and impatient firms). *Firms' discount factors are uniformly distributed on the open unit interval, $\delta_F \sim \mathbf{U}(0,1)$. A firm's respective δ_F is its private information.*[20] *Instead of equation (2.5), henceforth it is only required*

[19]Note that here we discount wages paid in a match period to the period preceding the match, as the worker incurs k_W and makes her decision to engage in search in this period.

[20]Note that none of Diamond's (1971) original assumptions is violated if we let workers be imper-

that $\exists\ e > 0$ *such that*

$$\frac{\delta_W[p(e) - c(e)]}{1 - \delta_W} > k_W + k_F \tag{2.13}$$

Now there are also firms that are too impatient to pay $\underline{b}$. (Throughout this section, we refer to those firms that deviate at a given b simply as impatient firms, and to the other firms as patient firms.) If these impatient firms truthfully advertised a $\tilde{b} < \underline{b}$, they would not attract any applicants. If they advertise some $\tilde{b} > \underline{b}$ but then renege already in their choice of $\hat{w}_1$, they will obtain a payoff 0. However, they can advertise some $\tilde{b} > \underline{b}$, thereby attract applicants, then pretend to play the SPE with $b_t = \tilde{b}$ by choosing $\hat{w}_1$ accordingly, and ultimately pay nothing after the worker exerted effort $e_1 = \bar{e}$. From then on, the SPE in lemma 2.2 would result, but such firms would already have obtained strictly positive payoffs this way.

A worker will thus only obtain a payoff $-c(\bar{e})$ in the first period of a match with an impatient firm. With δ_F^* here referring to the δ_F of the firm that is indifferent between deviating and paying $\underline{b}$, the worker's expected payoff is

$$(1 - \delta_F^*)\frac{\delta_W \underline{b}}{1 - \delta_W} - \delta_F^* \delta_W c(\bar{e}) < \frac{\delta_W \underline{b}}{1 - \delta_W} \tag{2.14}$$

where $1 - \delta_F^*$ and δ_F^* are the probabilities of meeting a patient and an impatient firm, respectively. Hence workers do not recoup their search costs in expectation and therefore do not engage in search.

However, a more desirable outcome might be achieved in a situation where patient firms manage to differentiate themselves from impatient firms. We will show that patient firms can indeed differentiate themselves if they sign a legally enforceable contract that specifies the time-invariant wage β and prescribes a severance payment σ in case the match is terminated. It is de facto irrelevant for the severance payment who terminates the match: even if it was formally prescribed only for firm-initiated terminations, workers wishing to terminate the match can in practice make the situation sufficiently unbearable to induce a firm-initiated termination. We clarify below how and when firms and workers in our set-up would choose to terminate the match.

fectly informed about firms' characteristics. Diamond's set-up only relies on firms having perfect information about workers' characteristics.

After the contract has been signed, it can only be changed by mutual consent. Hence a firm that signs the contract (β, σ) then must pay β in each period as long as the match lasts; if the match is terminated in some period, the firm must pay σ once in the same period. The contract may be signed in the interview in period 1^0 upon acceptance of the match by the worker and thus before 1^1 when the worker chooses effort. Histories can now be expressed as

$$\begin{aligned} h(1^0) &= \{\tilde{b}\} \\ h(1^1) &= \{\tilde{b}, \beta, \sigma, \hat{w}_1\} \\ h(t^0) &= \{\tilde{b}, \beta, \sigma, (\hat{w}_1, e_1, w_1), (\hat{w}_2, e_2, w_2), \ldots, (\hat{w}_{t-1}, e_{t-1}, w_{t-1})\} \end{aligned}$$

Without a contract $(\beta = \sigma = 0)$, the firm freely chooses the actual wage payments as before. We drop the time subscripts here because we found in the previous sections that b_t is time-invariant in all SPE of interest. If a contract is signed and wage β is paid, the definition of b as the part of the wage over and above $c(e)$ implies $b = \beta - c(e)$. Provided the severance payment is sufficiently high, only firms that are sufficiently patient to pay β sign the contract: while σ is not a cost to patient firms, it would have to be paid by impatient firms and thus makes deviations less attractive. Firms that do not sign a contract with a sufficiently high σ can then be identified as impatient firms. The severance payment functions here as a costly signal of being patient, just as education is a costly signal of having high ability in Spence (1973).

Apart from allowing a worker to identify patient firms, the contract may also affect workers' behaviour in more subtle ways: they might prefer to just receive the severance payment without working. That is, also workers might now be too impatient: the present discounted value to them of receiving b forever might be less than σ. If so, the worker will terminate the match in the first period immediately after the contract is signed. (Then the worker also does not work in the first period, as this would only create costs $c(e)$ for her because the firm has no reason to pay the wage in a terminating match.) For workers not to terminate the match, the following is needed:[21]

$$\sigma - \frac{k_W}{\delta_W} \leq \frac{b}{1-\delta_W} - \frac{k_W}{\delta_W} \tag{2.15}$$

If $\underline{b}$ satisfies this condition, profit-maximising firms will not pay more than that (so that $\beta = c(e) + \underline{b}$). However, if equation (2.15) is satisfied only at a $b > \underline{b}$, then the necessity to pay a sufficiently high σ will imply a minimum for b that exceeds $\underline{b}$. Firms will nevertheless pay this minimum for b to avoid termination of the match (but profit-maximising firms would not pay more than this). Importantly, this logic would explain *why firms pay more than needed to reimburse workers for search costs*: the firm's need to signal its patience gives rise to σ and then potentially to some $b > \underline{b}$.[22]

Patient firms can only differentiate themselves through σ if impatient firms do not want to mimic patient firms. In other words, there must not be any incentive in equilibrium for an impatient firm to engage in search, sign the same contract (β, σ) as patient firms, and then terminate the match. In particular, this needs to hold for the firm with the lowest δ_F that still engages in search and signs such a contract. From equation (2.15), we can conclude that patient firms seek to keep σ as low as possible in order to keep b as low as possible while still ensuring that the worker does not terminate the match. Hence the firm with the lowest δ_F that still engages in search will be indifferent between deviating and not deviating; otherwise the patient firms could reduce σ without any consequences. This firm's δ_F is therefore δ_F^*:

$$p(e) - \frac{k_F}{\delta_F^*} - \sigma = \frac{p(e) - c(e) - b}{1-\delta_F^*} - \frac{k_F}{\delta_F^*} \tag{2.16}$$

The left-hand side is the payoff from deviating in the first match period. Given that the firm is in the same situation in each match period, it deviates in the first period if at all. The firm deviates by not paying β at the end of the first period, which verifiably violates the contract and thereby terminates the match. The firm then has

[21] In equations (2.15) through (2.17), we discount to the first match period for simplicity.

[22] Beaudry (1994) uses a signalling game in a similar way to explain why, in a principal-agent context with effort and asymmetric information about job characteristics, even a perfectly informed principal might leave an agent with rents.

to pay σ but also appropriates $p(e)$ because the worker has already produced. The right-hand side is the payoff from not deviating: the firm obtains $p(e) - \beta$ in each period forever.

All firms with $\delta_F > \delta_F^*$ strictly prefer not to deviate. In any equilibrium where patient firms differentiate themselves by signing a contract (β, σ) with a sufficiently high σ, a firm that does not sign such a contract is treated as an impatient firm, thus fails to make any worker exert effort, and prefers not to engage in search so as to avoid k_F. Since the firm with $\delta_F = \delta_F^*$ is the firm with the lowest δ_F that still engages in search and pays σ, it must be indifferent between doing so and not engaging in search:

$$\frac{p(e) - c(e) - b}{1 - \delta_F^*} - \frac{k_F}{\delta_F^*} = 0 \tag{2.17}$$

Equation (2.16) implicitly assumed that the worker chooses $e > 0$. While the worker does not want to terminate the match because b has been chosen to satisfy equation (2.15), the worker might stop working without terminating the match, given that $p(e)$ is not verifiable. If the firm then preferred not to terminate the match either, the worker would receive β forever without working. However, σ is chosen such that the equality in equation (2.16) holds with $p(e) > 0$. Since $p(e)/(1 - \delta_F^*)$ exceeds $p(e)$, we can infer that

$$\frac{c(e) + b}{1 - \delta_F^*} > \sigma$$

is implied by the equality in equation (2.16). Hence the firm with δ_F^* strictly prefers to terminate the match and pay σ rather than pay β forever when $e = 0$. More patient firms have $\delta_F > \delta_F^*$, so that this preference is not reversed. Therefore, all patient firms terminate the match when the worker does not work, which by equation (2.15) is undesirable for the worker. The worker thus indeed chooses $e > 0$.

An equilibrium in which patient firms differentiate themselves from impatient firms would have to satisfy equations (2.15) through (2.17) together with equation (2.8). Equilibrium values, however, differ depending on whether or not $\underline{b}$ satisfies equation (2.15). Let us first suppose that equation (2.15) is not satisfied at $\underline{b}$. The following proposition claims that an equilibrium exists for a range of parameter values.

We analyse this signalling equilibrium as a PBE, which requires beliefs: let $\nu(\delta_F \geq \delta_F^*|\sigma)$ be a worker's subjective probability that a firm is patient, given that it signs a contract with severance payment σ.

Proposition 2.3 (PBE with sign-up bonus and $b^* > \underline{b}$). *$\exists$ a PBE of the model with $\delta_F \sim \mathbf{U}(0,1)$ in which*

(i) each firm that engages in search signs a contract with severance payment

$$\sigma^* = \frac{c(e^*) - k_F}{\delta_W}$$

(ii) workers' beliefs about δ_F are $\nu(\delta_F \geq \delta_F^*|\sigma \geq \sigma^*) = 1$ and $\nu(\delta_F \geq \delta_F^*|\sigma < \sigma^*) = 0$

(iii) workers and only those firms with $\delta_F \geq \delta_F^*$ engage in search, where

$$\delta_F^* = \frac{\delta_W k_F}{\delta_W p(e^*) - c(e^*) + k_F}$$

(iv) all firms that engage in search advertise and pay

$$b^* = \frac{1 - \delta_W}{\delta_W} [c(e^*) - k_F] > \underline{b}$$

(v) workers trust all potentially truthful advertisements:
$\mu(b = \tilde{b}|0 \leq \tilde{b} < p(\bar{e}) - c(\bar{e})) = 1$ and $\mu(b = \tilde{b}|\tilde{b} \geq p(\bar{e}) - c(\bar{e})) = 0$

(vi) the SPE profile $(\hat{w}_t, e_t, w_t) = (c(e^*) + b^*, e^*, c(e^*) + b^*)$ is played, $\forall t$ of a match where firms demand a level $e^* < \bar{e}$ that satisfies (primes denoting derivatives)

$$\delta_W p'(e^*) = c'(e^*)$$

provided that parameter values satisfy $k_W < c(e^*) - k_F < \delta_W[p(e^*) - k_F]$.

Proof. We find the equilibrium values by solving equations (2.15) through (2.17) and (2.8) simultaneously. Note first that equations (2.16) and (2.17) together imply

$$p(e) - \frac{k_F}{\delta_F^*} - \sigma = 0$$

We solve for δ_F^* as

$$\delta_F^* = \frac{k_F}{p(e) - \sigma} \tag{2.18}$$

We next rearrange equation (2.17) as

$$p(e) - c(e) - b = k_F \frac{1 - \delta_F^*}{\delta_F^*} \tag{2.19}$$

Substituting for δ_F^* from equation (2.18) leads to

$$\begin{aligned} p(e) - c(e) - b &= k_F \frac{p(e) - \sigma - k_F}{k_F} \\ \Leftrightarrow \quad \sigma^* &= c(e) - k_F + b \end{aligned} \tag{2.20}$$

With $b = (1 - \delta_W)\sigma$ as implied by equation (2.15) we obtain

$$\sigma^* = \frac{c(e) - k_F}{\delta_W} \quad \text{and consequently} \quad b^* = \frac{1 - \delta_W}{\delta_W} [c(e) - k_F] \tag{2.21}$$

This b^* exceeds $\underline{b}$ whenever

$$\begin{aligned} \frac{1 - \delta_W}{\delta_W} [c(e) - k_F] &> \underline{b} = \frac{1 - \delta_W}{\delta_W} k_W \\ \Leftrightarrow \quad c(e) - k_F &> k_W \end{aligned} \tag{2.22}$$

Workers will thus correctly believe to face a patient firm if and only if the contracted severance payment is at least σ^* and will exert effort only in this case. They correctly believe that they can trust advertisements including those $\tilde{b}$ arbitrarily close (but not equal or greater) to $p(\bar{e}) - c(\bar{e})$: by lemma 2.3, such SPE might be played by firms with sufficiently high δ_F. Profit-maximising firms, however, will pay the lowest possible b, i.e. b^* because anything lower would violate equation (2.15), so that workers would terminate the match. All firms thus pay $b^* > \underline{b}$, so that workers engage in search and match at the first opportunity (as do firms). Next, using the result for σ^* in equation (2.18) gives us

$$\delta_F^* = \frac{\delta_W k_F}{\delta_W p(e) - c(e) + k_F} \tag{2.23}$$

We require that $\delta_F^* < 1$, or

$$\delta_W k_F < \delta_W p(e) - c(e) + k_F \tag{2.24}$$

Because $\delta_W k_F > 0$ this condition, if satisfied, will also imply $\delta_F^* > 0$. The requirements on parameter values in equations (2.22) and (2.24) combine to

$$k_W < c(e) - k_F < \delta_W [p(e) - k_F]$$

Finally, the level of e demanded by firms satisfies equation (2.8). Employing the result for b^*, we rewrite this equation as

$$e^* = \arg\max_e p(e) - \frac{1}{\delta_W} c(e) + \frac{1-\delta_W}{\delta_W} k_F$$

Then the first-order condition with respect to e leads to

$$\delta_W p'(e^*) = c'(e^*)$$

Recall that $\bar{e}$ would equate $c'(e)$ with $p'(e)$. Since $\delta_W < 1$, e^* equates $c'(e)$ with something lower. By the convexity of $c(e)$, this equality must occur at some $e < \bar{e}$. □

The equilibrium in proposition 2.2 thus has the distinctive feature that firms pay more than needed to reimburse workers. In the presence of sufficiently high severance payments, workers would otherwise terminate the match. It appears that this logic in and of itself proposes a different solution to the Diamond paradox, without the need for advertisements. This solution would, however, be limited to certain parameter values, while we show below that firms will pay at least $\underline{b}$ also for other parameter values, thanks to advertisements. Finally, the last result is also interesting in its own right: here firms have an incentive to demand a lower level of e so as to limit b^*. This creates an inefficiency because $\bar{e}$ would maximise the surplus from match production.

Now suppose that $\underline{b}$ does satisfy equation (2.15). The logic of proposition 2.1 implies that firms will then advertise and pay $\underline{b}$ instead of a lower b that just satisfies equation (2.15), so that equation (2.15) can be slack in equilibrium. The equilibrium in this case has the same structure as before but different values:

Proposition 2.4 (PBE with sign-up bonus and $\underline{b}$). *$\exists$ a PBE of the model with $\delta_F \sim \mathbf{U}(0,1)$ in which*

(i) each firm that engages in search signs a contract with severance payment

$$\sigma^* = c(e) - k_F + \underline{b}$$

(ii) workers' beliefs about δ_F are $\nu(\delta_F \geq \delta_F^* | \sigma \geq \sigma^*) = 1$ and $\nu(\delta_F \geq \delta_F^* | \sigma < \sigma^*) = 0$

(iii) workers and only those firms with $\delta_F \geq \delta_F^*$ engage in search, where

$$\delta_F^* = \frac{k_F}{p(e) - c(e) + k_F - \underline{b}}$$

(iv) all firms that engage in search advertise and pay $\underline{b}$

(v) workers trust all potentially truthful advertisements:
$\mu(b = \tilde{b} | 0 \leq \tilde{b} < p(\bar{e}) - c(\bar{e})) = 1$ and $\mu(b = \tilde{b} | \tilde{b} \geq p(\bar{e}) - c(\bar{e})) = 0$

(vi) the SPE profile $(\hat{w}_t, e_t, w_t) = (c(\bar{e}) + \underline{b}, \bar{e}, c(\bar{e}) + \underline{b})$ is played, $\forall t$ of a match

provided that parameter values satisfy $-\frac{1-\delta_W}{\delta_W} k_W \leq c(e) - k_F \leq k_W$.

Proof. Equation (2.20) can be derived exactly as in the proof of proposition 2.2, with the only difference that $b = \underline{b}$ is already determined. Here therefore

$$\sigma^* = c(e) - k_F + \underline{b} \tag{2.25}$$

If the severance payment is at least σ^*, workers will again correctly believe to face a patient firm. Plugging the result for σ^* back into equation (2.18) gives

$$\delta_F^* = \frac{k_F}{p(e) - c(e) + k_F - \underline{b}} \tag{2.26}$$

Since, by assumption 2.4',

$$p(e) - c(e) > \frac{1 - \delta_W}{\delta_W} k_W = \underline{b} \tag{2.27}$$

we conclude that $\delta_F^* \in (0, 1)$ for any parameter values. However, we require $\sigma^* \geq 0$ and recall that the PBE in proposition 2.2 applies whenever equation (2.22) holds. Hence the PBE in proposition 2.3 will only exist if parameter values satisfy

$$-\frac{1 - \delta_W}{\delta_W} k_W \leq c(e) - k_F \leq k_W \tag{2.28}$$

That workers correctly believe they can even trust $\tilde{b}$ arbitrarily close to $p(\bar{e}) - c(\bar{e})$ follows from lemma 2.3 and the proof of proposition 2.1 as before. Also by the same logic as before, workers engage in search and match at the first opportunity. During the match, firms demand the level of e determined by equation (2.8). As $\underline{b}$ does not depend on e, this level is $\bar{e}$ here. $\square$

The overall intuition for propositions 2.2 and 2.3 is that the firm has to gain the worker's trust by contracting for a severance payment σ, which the firm will only do if it is patient. What sustains these PBE is the loss of trust in the firm that would occur if the severance payment was less than σ^*. The core logic here is thus a very simple signalling game. Propositions 2.2 and 2.3 together show for a connected range of parameter values that, when there are impatient firms, a severance payment can restore the result that workers engage in search. While contracts with severance payments *can* separate patient from impatient firms, the question remains whether this necessarily happens whenever such contracts exist. The following corollary provides an affirmative answer:

Corollary 2.1 (Uniqueness). *Suppose we require that beliefs are consistent with probability 0 of strictly dominated and equilibrium-dominated actions occurring. Then* σ^* *in proposition 2.2 is the unique* σ *in equilibrium for any equilibrium in the range of parameter values for which proposition 2.2 holds, and likewise for* σ^* *in proposition 2.3.*

Proof. We only sketch the proof here because structurally identical proofs have been provided in similar contexts. The following argument applies equally to propositions 2.2 and 2.3. Essentially, any $\sigma > \sigma^*$ is a strictly dominated action for an impatient firm, so that the worker should believe to face a patient firm when any such σ is contracted for. Hence patient firms need not contract for more than σ^* to credibly signal their patience. Then only the separating equilibrium with σ^* remains among the separating equilibria. Patient firms can and will break every pooling equilibrium by choosing a σ that would be an equilibrium-dominated action if chosen by an impatient firm, so that workers identify such deviants as patient firms. $\square$

Under the conditions on beliefs in corollary 2.1, contracts with severance payments thus necessarily separate patient firms from impatient firms. Hence, the existence of such contracts effectively plays the same role as the assumption of patient firms we had made before proposition 2.1. The equilibria in propositions 2.2 and 2.3 thus represent the PBE in proposition 2.1 extended by a signalling game instead of assumption 2.4. In conclusion, it is not necessary that all firms be patient to arrive at the finding that there are equilibria in which workers engage in search.

This section has thus explained why impatient firms undermine the PBE in proposition 2.1 and has argued that contracts with severance payments necessarily fix this problem for a connected range of parameter values. For the part of this range where $k_W < c(e^*) - k_F$, proposition 2.2 establishes in addition that all firms advertise and pay a b^* strictly above $\underline{b}$, so that workers earn strictly positive rents. By our assumptions 2.1 and 2.2, none of these results rely on competition between firms or on imperfectly informed firms.

2.7 Conclusions

We hope this chapter furthers our understanding why forward-looking job seekers participate in labour markets at all. Our analysis suggests that advertisements and efficiency wages are both important parts of an explanation. Including these elements in a model with search frictions and perfectly informed firms that set wages, we have found equilibria in which workers participate and may even earn rents. In these equilibria, all firms pay the same wage. Contrary to what is sometimes believed, solving the Diamond paradox is not necessarily the same as generating wage dispersion. Unlike responses to the Diamond paradox in the literature, the solution proposed here does not somehow rely on imperfectly informed firms or their ability to commit. By relying on advertising instead, our solution rather emphasises the role of labour market institutions. While binding minimum wages or unions have also been identified as solutions to the Diamond paradox, these latter institutions do not appear as widespread in real-world labour markets as advertising.

Chapter 3

Assortative matching through signals

Imagine a dark room filled with a number of married couples. The people in the room are busy attempting to find their mates. An economist would attempt to model the behavior in the room by assuming that everyone is searching for his better half in an optimal manner given the fact that the lights are out. Consequently, he would conclude that the behavior in the room is efficient. However, such a conclusion does not deny that the matching process would be improved if someone should happen to turn on the lights.

Dale T. Mortensen (1976), p. 195, echoing a parable told by Axel Leijonhufvud.

3.1 Introduction

A number of important markets, notably the labour market, bring trade partners together in pairs that form durable matches. The dominant approach in the analysis of decentralised interaction on such matching markets has been developed by Diamond, Mortensen, and Pissarides: agents engage in search for match partners, and frictions in the search process make this costly. While this approach and derived results have found ample empirical support, efforts to replicate important empirical regularities in search models continue. Alongside business cycle properties and price or wage dispersion, it is the pattern of sorting in matching markets that the literature has sought to explain. Indeed, that likes tend to match with likes is a pervasive phenomenon: more productive workers tend to be hired by more productive firms, more educated women tend to marry more educated men, more reliable tenants tend to secure nicer apartments.[1]

This sorting of likes along some dimensions, known as *positive assortative matching* (PAM), thus serves as a test on the validity of the search model: whenever a model of a decentralised matching market with heterogeneous agents cannot generate PAM, it appears to have missed something important. That apart, PAM is relevant because it may be more efficient than, say, random matching: highly skilled workers are better placed in complex work environments than unskilled workers, for example. If economists understand what drives assortative matching, they will be in a position to avoid mismatch in market design or to improve efficiency in existing markets by making PAM more pronounced.

The theoretical literature has identified some form or another of complementarity among the inputs into a match as the driver of PAM. Thinking of the value generated by a match as the match output, a match production function specifies how agents' inputs translate into match output. Then simple complementarity among the inputs is equivalent to *supermodularity* of the match production function: the marginal effect on output from a change in one input is increasing in the other input. A stronger concept of complementarity, log-supermodularity, requires that the marginal effect on output from a change in one input is also as a proportion increasing in the other input.

While supermodularity as such suffices for PAM to arise in the model whenever search is costless, it often does not suffice when search is costly, in particular when search takes valuable time, which it does in practice. The literature has therefore resorted to more demanding forms of complementarity, but it is doubtful that very

[1] As an exemplary reference for these stylised facts, see Mare (1991).

demanding conditions hold across the numerous and diverse real matching markets that exhibit PAM. In this chapter, however, we find that a model with a greater availability of information in the search process, equivalent to more light in the quote from Mortensen above, only requires substantially weaker conditions to generate PAM: search with better information takes less time and avoids some costly but unsuccessful meetings, so that our model is in effect closer to models with costless search, in which supermodularity as such is sufficient. This chapter therefore reconciles the search-theoretic approach with potentially very many decentralised markets that exhibit PAM in practice but do not meet the existing theoretical conditions for PAM.

State of the literature Following Smith (2006), we would classify the models in the literature by two criteria. One is whether or not utility is transferable between agents. The literature refers to *non-transferable utility* whenever agents divide the match output according to a pre-imposed split, like an employer and her employee do when the employee's wage is already determined by a union wage agreement. In cases of *transferable utility*, there is no pre-imposed split and agents have to bargain over the match output and agree on a split. The other criterion is the kind of frictions in the model as captured by the costs of search. The seminal article by Becker (1973) considers a frictionless setting (i.e. agents search costlessly) where utility is transferable, and supermodularity as such turns out to suffice for PAM, in fact even perfect PAM: only equal types match.[2]

In a setting with frictions, the costs of search may be expressed as implicit costs through discounting or as explicit additive costs. An influential contribution by Shimer and Smith (2000) examines a setting with discounting and transferable utility. They conclude that the match production function, the logarithm of its first derivative, and the logarithm of its cross-partial derivative all need to be supermodular for PAM to arise in this setting. For the same setting with non-transferable utility, Smith (2006) shows that PAM will only arise if the match production function is log-supermodular. While supermodularity as such is a very natural condition,

[2]Using Becker's (1973) work, it is quickly found that a frictionless setting with non-transferable utility leads to PAM even without supermodularity: see Smith (2006), section II.

the combination of three conditions required for PAM in Shimer and Smith (2000) is criticised by Atakan (2006) as "quite restrictive"; Goldmanis et al. (2009) find it "quite troubling" that supermodularity is not sufficient in Smith (2006) and point out that even situations in which there is hardly any or no sorting at all meet the formal criterion for PAM that these papers employ.[3] In short, there is a paradox here: real-world agents in all sorts of matching markets both discount and sort into PAM, but theoretical models of these markets require improbably strong complementarities to generate even a very weak form of PAM.

Atakan's (2006) own model features transferable utility, as in Shimer and Smith (2000), but explicit costs instead of discounting. Morgan (1998) builds a model with non-transferable utility and explicit costs. In both models, supermodularity as such gives rise to PAM (while the results in Atakan (2006) crucially depend on search costs being identical for all agents). However, limiting oneself to explicit costs has not provided an answer to the paradox, since real-world agents do discount.

This chapter We offer a solution to this paradox by appealing to another, equally pervasive phenomenon in matching markets: the use of signals. Signals come in the form of job advertisements and applications on the labour market, dress style and body language on the marriage market, online photographs of apartments and flat hunters on the market for rented housing, to name but a few. We argue that such signals, if only they are truthful, make all wasteful search obsolete: agents can target their search directly on the best offer they can possibly secure, and they thus never lose money or time on search that they regret. By cutting search short, signals can reduce the effect of search frictions and thus bring a setting with frictions closer to a frictionless setting, in which mild conditions suffice for PAM to arise.

Of course, signals will only have this beneficial effect if they are actually informative. In the famous signalling model by Spence (1973), years of education are an informative signal for ability because more able workers find it easier to acquire education than less able workers do. The model thus relies on signals being costly and

[3]See Atakan (2006), p. 667 and Goldmanis et al. (2009), p. 8 and p. 12.

on a single-crossing property. Unfortunately, when agents signal their type through applications or advertisements, the costs are typically small and a single-crossing property can in general not be expected to hold: writing a forged CV is as costly as writing a truthful CV, and painting an advertised job in unduly bright colours is as costly as honestly laying out its dull nature. Yet, as shown by contributions such as Crawford and Sobel (1982), signals will be informative even in an environment of such *cheap talk* if the interests of senders and receivers are sufficiently aligned. Then the costs of signals can just as well be normalised to zero.

In essence, the model we build introduces costless signals into the search model with transferable utility in Shimer and Smith (2000). We prove a unique separating equilibrium in which matching is not only positively assortative but perfectly positively assortative if and only if the match production function is supermodular. In this equilibrium, each agent finds it optimal to signal truthfully and to target her search on only one type. A low type will in fact not prefer to deviate to matches with higher types in equilibrium because the higher types' optimal behaviour in effect forces her not to renege on her signal. When reneging is not an option, the lower type has to conceal the difference between expected and actual match production by reducing her share accordingly. If the match production is supermodular, this reduction will outweigh the gain from matching with a higher type. Intuitively, with truthful signals perfect PAM is then to be expected: if signals are fully informative, agents can replace (almost all) costly search via meetings by costless search via signals and can therefore behave like in a frictionless setting.

The separating equilibrium has a number of desirable efficiency properties. Above all, agents match at the very first opportunity, so that no time is wasted on unsuccessful search and search costs are minimised. In labour market terms, this would mean that frictional unemployment is reduced to a minimum. Instead of the conditions identified in Shimer and Smith (2000), just the weaker condition of supermodularity suffices in our model to generate even perfect PAM instead of merely PAM. This chapter thereby offers a resolution to the paradox created by the conditions in Shimer and Smith (2000). Finally, to the best of our knowledge, ours is the only model that generates perfect PAM despite discounting, or explicit search costs, or both.

Relation to the literature Only few papers consider sorting in the context of a matching market with signals. Hoppe et al. (2009) and Hopkins (2010) build two similar models of a *matching tournament* with signalling: match partners are essentially prizes for ex-ante choices of costly signals. In both models, agents first select a costly signal of their unobservable type like in Spence (1973) and then, based on these signals, match roughly like in Becker (1973). Hopkins (2010) assumes a single-crossing property and Hoppe et al. (2009) assume a specific multiplicative production that satisfies supermodularity and even log-supermodularity. In the symmetric equilibrium, agents' signals are then strictly increasing in their types. This leads to perfect PAM at the matching stage - just as one would have expected, given Becker's (1973) findings. However, since search frictions do not exist in matching tournaments, neither of the two papers helps us to resolve the paradox in Shimer and Smith (2000).

A search model built by Chade (2006) features discounting and noisy signals uncontrolled by the agents. Yet these signals are not observed before agents meet. Rather, when agents do meet, they do not observe each others' true types but only the noisy signal. Hence search is still random in this model, and the noisy signals in fact add information frictions to the search frictions. Assuming that the noisy signals exogenously carry some information, matching is shown to exhibit PAM in a very weak sense: the distribution of types a high type might match with first-order stochastically dominates this distribution for a low type. Chade (2006) further finds that an assumption of log-supermodularity of the match production function, while not necessary, reinforces these results.

Our set-up primarily differs from Chade's (2006) in that signals are observed before meetings and thus make search non-random, thereby tending to reduce search frictions. Moreover, signals in our model are not exogenous truthful signals with noise but are deliberately and strategically chosen by agents. We would argue that real-world agents will exert as much control as possible over the signals of their types, given how important the signals can be for their payoffs. Hence, we do not assume that signals are necessarily informative.[4]

Further, a model by Eeckhout and Kircher (2010) features the key elements of models of *directed search*, including signals in the form of sellers' posted offers. Buyers observe the posted offers and simultaneously choose which seller to visit. The only way frictions enter this process is through congestion: buyers cannot coordinate, so that queues result and only some buyers manage to buy. It is shown that, for common meeting technologies, the matching of buyers and sellers will exhibit PAM if the square root of the match production function is supermodular. Shimer (2005a) looks at PAM in a directed search model of the labour market and finds, at least for the case of only two worker types, that there will be some stochastic form of PAM as long as workers of low type do not have a comparative advantage when working for employers of high type. While these models simply assume that signals are truthful (a criticism we would level against almost any directed search model), these models cannot resolve the paradox mainly because of their limitation to frictions from congestion, so that their results do not carry over to models with other frictions in any obvious way.

A recent contribution by Lentz (2010) does not feature any signals but allows for search on the job (or, more generally, search while matched) with endogenous search intensity in a model with discounting. While search is still random, higher types gain more than others from search on the job if the match production function is supermodular. Higher types thus search more intensively, which leads to PAM in terms of stochastic dominance as in Chade (2006). A related model by Goldmanis et al. (2009) features non-transferable utility, discounting, and search on the job. If only one agent in each match can switch to another match, PAM will result if the match output function is log-supermodular. If either agent can switch, a condition has to be met so that agents gain from switching to matches with higher types (at least in the equilibrium they identify, which might not be unique). Then the situation would

[4] Visschers (2005) makes another contribution that adds information frictions to the search frictions: upon meeting, agents only observe each other's signals, which are costly and strategically chosen. He argues that, with non-transferable utility and log-supermodularity as in Smith (2006), such signals can lead to block segregation rather than PAM, in contrast with the results in Smith (2006).

evolve into perfect PAM over time, were it not always set back as matches randomly break up and agents begin from scratch. Therefore, perfect PAM is only achieved in the limit as the rate of random break-ups tends to zero. Moreover, agents in Lentz (2010) and Goldmanis et al. (2009) sort only over time. The fundamentally different sorting mechanism in our model can, by contrast, explain sorting already among graduates in their first job and achieves perfect PAM in the unique (separating) equilibrium of a model with random break-ups, without invoking stronger conditions.

In a subtle way, our model is finally also related to Jacquet and Tan (2007). They consider an environment with discounting and a particular log-supermodular match production function: an agent's payoff equals her match partner's type (so that utility is non-transferable). For such an environment, Burdett and Coles (1997) found that types segregate into classes and match exclusively within these classes. Building on this, Jacquet and Tan (2007) let agents establish any number of marketplaces and show that each marketplace is populated by only one class in equilibrium. By going to the appropriate marketplace, each agent can thus avoid meetings that do not lead to a match and can instead match after the first meeting. Agents can do the same in our model if signals are informative. In a way, the signals in our model can create any number of virtual marketplaces. While it is left open in Jacquet and Tan (2007) how marketplaces are established and located by agents, virtual marketplaces are established and accessed simply by selecting among the observed signals. Apart from this difference, our environment will feature transferable utility and a general match production function.

Overview The chapter proceeds as follows. Section 3.2 specifies a frictional matching market and the procedures of search. Section 3.3 defines equilibrium in the model and proposes a separating equilibrium in which supermodularity suffices for perfect PAM. Its existence is proven step by step through a series of lemmas in section 3.4. The separating equilibrium is found to be unique as well as efficient in section 3.5 before section 3.6 concludes.

3.2 Model

The matching market in our model consists of heterogeneous agents who match among themselves. Agents are indexed by a productivity type $x \in \Theta$, where $\Theta = [\underline{x}, \bar{x}] \subset \mathbb{R}_{++}$. For each type, there is a continuum of agents. The measure of agents with types weakly below $x \in \Theta$ is denoted $L(x)$ and assumed C^1. We require its density function $l(\cdot)$ to be everywhere strictly positive. After normalising the overall mass of agents to one, $l(x)$ may be thought of as the mass of agents of type x. We denote the density function for unmatched agents by $u(\cdot) \leq l(\cdot)$. Then $u(x)$ represents the mass of unmatched agents of type x, while $l(x) - u(x)$ represents the mass of matched agents of type x. These quantities are determined endogenously.

Types are exogenously given, but only privately observable. To convey information about their type, agents can choose any signal $\tilde{x} \in \Theta$ as a signal of their type. To distinguish in the notation between two agents who might match, we will denote the type of one agent by x and the type of the other by y. Normalising also the flow output generated by an unmatched agent to zero, a match between types x and y generates a constant flow output $f(x, y)$. The match production function $f(\cdot, \cdot)$ is assumed to satisfy some regularity conditions:

Assumption 3.1 (Regularity conditions). *The match production function $f(\cdot, \cdot)$ is symmetric ($f(x, y) \equiv f(y, x)$), increasing in both arguments, and, for any existing types, takes only positive values ($f : \Theta^2 \mapsto \mathbb{R}_{++}$).*

Because productivity types are scalars, there can only be gains from specialisation in match production if one role in production rewards productivity more than the other role does. While the assumption of symmetry of $f(\cdot, \cdot)$ does not rule out such specialisation, it amounts to the assumption that the more productive agent always assumes the role that rewards productivity more, so that match output is maximised.

Time is continuous with an infinite horizon. Each agent is always in one of three states: matched, searching, or not participating. When indifferent whether to engage in search and whether to accept a match, an agent does respectively search and accept the match. An agent who engages in search pursues two types of activities that neither matched nor non-participating agents pursue. Firstly, she sends a costless and public signal of her type. For some agent x, this is $\tilde{x} \in \Theta$, which may or may not be an accurate signal of her true type x. She can always instantly and costlessly change the signal she is sending. Secondly, an active agent seeks meetings: before x can match with some agent y, a meeting between the two will have to occur. Given the existence of signals, agents can condition their willingness to meet on observing particular signals.

However, observing a signal $\tilde{y}$, agent x can only form a belief about the true type y behind the signal. Agents never directly observe each other's actual types. Let h^x be the history of one agent observed by x, i.e. a set of actions such as the observed signal. We represent a belief as a probability distribution $\Psi(\cdot)$. Concretely, for each h^x, the belief held by agent x of the other agent's true type y is the probability distribution $\Psi(\cdot|h^x)$ over Θ. Then x, having observed h^x, believes that the other's type is y with probability density $\psi(y|h^x)$. All agents use Bayes' rule to form and update their beliefs. Further, note that match output must also be unobservable when types are unobservable: knowing $f(\cdot, \cdot)$, x could otherwise recover y from the observed output $f(x, y)$. To keep the notation simple, let $g(x|h^x)$ denote the match output that x expects based on her belief after observing h^x:

$$g(x|h^x) = \int_{y=\underline{x}}^{y=\bar{x}} f(x, y)\psi(y|h^x)dy$$

Precisely because agents can condition meetings on signals, meetings are non-random. Agents who limit their meetings this way effectively create marketplaces where only those sending certain signals meet each other. Agents enter such a marketplace simply by sending one of the required signals and by seeking to meet only agents who also send one of the required signals. We assume that inside any such marketplace, meetings are random and described by a function $m(\cdot)$ with $m(0) = 0$. The flow of meetings in a marketplace with a mass of agents λ then equals $m(\lambda) \leq \lambda$. Assuming constant returns to scale in meeting, the meeting rate η is the same across all marketplaces:

$$\eta = \frac{m(\lambda)}{\lambda}, \quad \forall \lambda$$

Agents can create any marketplace that attracts some agents. Each agent chooses which marketplace to join by choosing whom to meet and which signal to send. Let $D(x)$ denote the set of signals $\tilde{y}$ such that an agent of type x would like to meet the agent sending such a signal, that is, the *meeting strategy* (or, in marriage market terms, the dating strategy) of agent x. Then the set of signals $\tilde{y}$ whose senders would be willing to meet x, given her signal $\tilde{x}$, is her *opportunity set*

$$\Omega(x) \equiv \{\tilde{y} : \tilde{x} \in D(y)\}.$$

Combining these sets, $R(x) \equiv D(x) \cap \Omega(x)$ denotes the set of signals $\tilde{y}$ such that a meeting (or a *rendez-vous*) between x and y would in fact result.

Since utility is transferable, agents in a meeting bargain over the division of the match output that they would produce between them. We model this using a strategic bargaining procedure with alternating offers. It is useful to imagine that $f(x, y)$ is contained in a pot that agents can only take from but cannot look into. (If $f(x, y)$ was observed, knowledge of $f(\cdot, \cdot)$ would allow each agent to infer the true type of the other agent.) When agents first meet, either of them is randomly selected with probability $\frac{1}{2}$ to move first. An agent x who moves first takes a share $\pi(x|y)$ for herself from the pot, which is not observed by y. Then y takes the remainder $f(x, y) - \pi(x|y)$ from the pot (which may be negative), unobserved by x. Now y has three options: she can accept the share left for her in the pot, reject this share, or walk away and immediately continue searching.[5] If y accepts, agents match immediately and obtain their respective share as a flow utility for the duration of the match. If y rejects, x can walk away; otherwise, the same two agents meet again for the next round of bargaining, in which one agent is again selected with probability $\frac{1}{2}$ to move first, and so on. Previously offered shares cannot be recalled, and if players never agree nor

[5]The fact that agents have met implies that these agents prefer engaging in search to not participating. It is thus without loss of generality that non-participation is not a further outside option here.

walk away, both will obtain 0.

In order to be more precise, let us formalise this bargaining game. The players are 'nature' N and the agents x and y who meet. The history h^x records the actions that x has observed thus far, h^y records what y observed, and we simply index histories in chronological order. When x and y first meet, they already know both signals, so that $h_1^x = h_1^y = \{\tilde{x}, \tilde{y}\}$. A player function $P(\cdot,\cdot)$ assigns to each history pair (h^x, h^y) (except any terminal history pairs) a player who moves at this history pair. $P(h_1^x, h_1^y) = N$ selects x and y each with probability $\frac{1}{2}$ to move first. If x is selected, then $h_2^x = h_2^y = h_1^x \cup \{x\}$ and $P(h_2^x, h_2^y) = x$. Agent x now chooses an action according to her bargaining strategy $B(x)$ that assigns an action to every possible history pair for which $P(h^x, h^y) = x$. If $P(h_2^x, h_2^y) = x$, x chooses some $\pi(x|y)$. As y observes only the remainder, $h_3^y = h_2^y \cup \{f(x,y) - \pi(x|y)\}$ while $h_3^x = h_2^x \cup \{\pi(x|y)\}$, and y responds according to $B(y)$ by choosing an action from the set {"accept", "reject", "walk away"}. If she chooses "accept" or "walk away", then (h_4^x, h_4^y) will be a terminal history pair. If she chooses "reject", then $h_4^x = h_3^x \cup$ {"reject"} and $P(h_4^x, h_4^y) = x$ who chooses from {"continue", "walk away"}. If x does not walk away, bargaining will continue in the next meeting with $P(h_5^x, h_5^y) = N$, and so on.

Assumption 3.2 (Common expected delay). *A further meeting with the same agent happens at the same rate as a new meeting with another agent.*

That is, a time $1/\eta$ elapses in expectation before another bargaining round, so that both a meeting to continue bargaining and a meeting with a different agent happen at rate η. While there is nothing in our model that would cause these meeting rates to differ systematically, assumption 3.2 is obviously a simplification. Agents who are currently waiting for another round of bargaining cannot engage in search at the same time, so that they do not send a signal while waiting. This avoids random meetings with other agents. Those between bargaining rounds therefore form a set of agents who cannot be encountered on any marketplace. The meeting rates on all marketplaces are unaffected due to constant returns to scale in meeting. Finally, matches dissolve exogenously at constant Poisson rate δ.

All agents are risk-neutral, discount future utility at discount rate r (with $0 < r < \infty$), and seek to maximise the present discounted value (pdv) of their expected utility. Throughout the chapter, 'payoff' refers to the pdv, not to the flow utility. Because of discounting, the time that elapses before a meeting makes meetings costly. In addition, we allow for explicit costs $c \geq 0$ that an agent incurs each time she attends a meeting.

Assumption 3.3 (Gains from trade). *The output produced in a match between two agents of the lowest type, discounted at effective discount rate $r+\delta$, can reimburse both agents' explicit costs of one meeting:*

$$2c \leq \frac{f(\underline{x}, \underline{y})}{r+\delta}$$

While explicit costs always have to be limited relative to the available payoffs to ensure agents' participation, note that this assumption is particularly mild. For example, we do not assume that each agent is in fact reimbursed in the event of a match, nor that match output is sufficient to reimburse the costs of the expected number of meetings before a match. With the inclusion of explicit costs in addition to discounting, our model thus features two kinds of search frictions. Finally, agents know everything except the true type of any other agent and, by consequence, the actual match output $f(x, y)$ they produce with another agent.

3.3 Equilibrium

3.3.1 Definition of equilibrium

Define $U(x)$ as the expected present value to x of being unmatched and engaging in search, and $W(x|y)$ as the expected present value to x from a match with y. Let the set $A(h^x, h^y)$ comprise of all combinations of bargaining strategies $(B(x), B(y))$ that lead to a SPE of the bargaining game given history pair (h^x, h^y), so that an agreement is reached in the meeting in question and agents match. Let us define an indicator function $\alpha(\cdot, \cdot)$ such that $\alpha(B(x), B(y)) = 1$ if $(B(x), B(y)) \in A(h^x, h^y)$

and 0 otherwise. Then the following *asset equation* expresses the expected return on being unmatched as the expected gain from a meeting net of search cost c:

$$rU(x) = \eta \left(-c + \int_{\Theta} \alpha(B(x), B(y)) \left[W(x|y) - U(x) \right] \psi(y|\tilde{y} \in R(x)) dy \right) \tag{3.1}$$

where $\psi(y|\tilde{y} \in R(x))$ is the density of y that x believes conditional on a meeting.

As is natural when signals are involved, we look for a *perfect Bayesian equilibrium* (PBE) of our model. We focus our attention on separating equilibria.[6] As signals are costless, all PBE will necessarily be cheap-talk equilibria. A steady-state PBE of our model, separating or not, requires that the flows into and out of matches balance for every type (a pointwise steady state), that all agents choose all their strategies optimally, and that agents' beliefs are consistent with all agents' actual equilibrium behaviour.

Definition 3.1 (Search equilibrium with signals). *In a steady-state PBE of our model, each agent $x \in \Theta$*

(i) engages in search if and only if $U(x) \geq 0$

(ii) chooses her signal optimally as $\arg\max_{\tilde{x}} rU(x)$ given $B(x), B(y), R(x)$ where $rU(x)$ is determined by equation (3.1)

(iii) chooses an optimal meeting strategy as $\arg\max_{D(x)} rU(x)$ given $B(x), B(y), \Omega(x)$, noting that $R(x)$ depends on $D(x)$

(iv) chooses a subgame-perfect bargaining strategy as $\arg\max_{B(x)} rU(x)$ given $B(y)$, $R(x)$, noting that $W(x|y)$ depends on $B(x)$

(v) holds beliefs that are formed using Bayes' rule where possible and are consistent with equilibrium play: given an equilibrium history h^x, $\psi(y|h^x) = l(y|h^x)$ where $l(y|h^x)$ is the true density of y conditional on history h^x

[6] Kübler et al. (2008) report experimental evidence suggesting that pooling equilibria never arise when some types can benefit from the effective use of signals.

and the matching market is in a pointwise steady state: given $B(x)$, $B(y)$, and $R(x)$, for all $x \in \Theta$

$$\delta[l(x) - u(x)] = \eta u(x) \int_{\Theta} \alpha(B(x), B(y)) l(y|\tilde{y} \in R(x)) dy \tag{3.2}$$

which also holds when x does not engage in search and thus $l(x) - u(x) = u(x) = 0$.

However, a PBE only requires agents' beliefs to be consistent with equilibrium play, not with actions out of equilibrium. As is well known, a PBE can therefore depend on unreasonable off-equilibrium beliefs because these beliefs are never tested in equilibrium. Since unreasonable beliefs are not needed for any of our results, we rule out beliefs that are unreasonable in the sense of the Intuitive Criterion. To do this formally, let us call the combination of $\tilde{x}$, $D(x)$, and $B(x)$ the 'grand strategy' of agent x, denoted $GS(x) = (\tilde{x}, D(x), B(x))$. Also define $BR(x|h^x)$ as the set of continuation strategies $GS(x|h^x)$ that are best responses for x. To apply the Intuitive Criterion as an equilibrium refinement, we have to define the notion of equilibrium domination in our model:

Definition 3.2 (Equilibrium domination). *Given a PBE of the model, the continuation strategy $GS(x|h^x)$ is equilibrium-dominated at history pair (h^x, h^y) if*

$$U^*(x) > \max_{GS(y|h^y) \in BR(y|h^y)} U(x|GS(x|h^x))$$

where $U^*(x)$ denotes the value of $U(x)$ in the PBE and $U(x|GS(x))$ is the value of $U(x)$ obtained with strategy $GS(x|h^x)$.

The Intuitive Criterion then demands that the beliefs of y about off-equilibrium actions place probability 0 on any x who would have to pursue equilibrium-dominated strategies to reach the off-equilibrium history: $\psi(x|h^y) = 0$ if, at a history up to h^y, x would have had to play an equilibrium-dominated strategy $GS(x|h^x)$.

3.3.2 Putative equilibrium

We next propose that a particular separating equilibrium exists under a simple condition on the match production function $f(\cdot,\cdot)$. The condition requires some complementarity among the inputs x and y. Concretely, all we need is a weak and intuitive form of complementarity known as strict supermodularity (or increasing differences): the marginal product of one agent in a match is strictly increasing in the type of the other agent.[7]

Definition 3.3 (Supermodularity). *The match production function $f(\cdot,\cdot)$ is strictly supermodular if, for all $x_H > x_L$ and $y_H > y_L$,*

$$f(x_H, y_H) - f(x_L, y_H) > f(x_H, y_L) - f(x_L, y_L)$$

Further, we refer to the sorting with $x = y$ in all matches as *perfect positive assortative matching* (PPAM). We can now propose existence of the following particular PBE in our model:

Proposition 3.1 (Existence). *Let agents' beliefs place probability 0 on the occurrence of equilibrium-dominated actions. Then for any type distribution $L(x)$, strict supermodularity of $f(\cdot,\cdot)$ is necessary and sufficient for the existence of a separating PBE in which each agent $x \in \Theta$*

(i) engages in search: $U(x) \geq 0$

(ii) signals truthfully: $\tilde{x} = x$

(iii) meets exclusively agents of her own type: $D(x) = \{\tilde{y} : \tilde{y} = x\}$

(iv) reaches a bargaining agreement in the first meeting and thus matches: $\alpha(B(x), B(y)) = 1$ for $x = y$

(v) correctly believes all signals to be truthful: $\psi(y|\tilde{y}) = l(y|\tilde{y}) = 1$ for all $y = \tilde{y}$

[7] A stronger form of complementarity is strict log-supermodularity, which is defined using $\ln f(\cdot,\cdot)$ instead of $f(\cdot,\cdot)$ in definition 3.3, so that the proportional marginal product of one agent is increasing in the other agent's type.

and the market is in pointwise steady state. The resulting equilibrium matching is PPAM.

A proof of proposition 3.1 will thus establish that not only PAM, but even PPAM arises in our model under the same weak condition as in a frictionless model, although our model allows for two kinds of frictions. In a model with frictions only from discounting, Shimer and Smith (2000) establish PAM, albeit not PPAM, under the condition that the match production function $f(\cdot,\cdot)$, the logarithm of its first derivative, and the logarithm of its cross-partial derivative are all supermodular. These conditions are directly comparable to our condition and are unambiguously more restrictive: proposition 3.1 claims that our model achieves PPAM with just the first of these conditions, which is also the most intuitive.

The next section proves proposition 3.1 through a series of lemmas. Each time, we separately consider a component of proposition 3.1, taking as given that all other components are indeed as specified in proposition 3.1. We verify for the component in question, as applicable, that it is optimal for agents to behave as specified in proposition 3.1, that a steady state results, or that beliefs place probability 0 on equilibrium-dominated actions and are consistent with equilibrium play.

3.4 Existence proof for the putative equilibrium

3.4.1 Bargaining

We first derive the expected present values that various states carry in the putative equilibrium situation. Given that beliefs are consistent with equilibrium play, we have

$$\psi(y|\tilde{y} \in R(x)) = l(y|\tilde{y} \in R(x))$$

If x only meets agents of her own type and matches with them, then

$$l(y|\tilde{y} \in R(x)) = \alpha(B(x), B(y)) = 0 \quad \forall y \neq x \tag{3.3}$$

Since every meeting in the putative equilibrium thus leads to match, the rate of matches equals the rate of meetings, and an agent effectively incurs costs c each time she matches. Therefore, equation(3.1) simplifies to

$$rU(x) = \eta\left[W(x|y) - c - U(x)\right] \tag{3.4}$$

with $y = x$. Next, the expected return on being matched with y is the expected flow utility and the loss from match dissolution at rate δ:

$$rW(x|y) = \sigma(x|y) - \delta[W(x|y) - U(x)] \tag{3.5}$$

where $\sigma(x|y)$ denotes the expected share that x obtains when bargaining with y over the flow of match output $f(x, y)$,

$$\sigma(x|y) = \frac{1}{2}\pi(x|y) + \frac{1}{2}[f(x, y) - \pi(y|x)] \tag{3.6}$$

One can solve equation (3.4) for $U(x)$ and equation (3.5) for $W(x|y)$, then use the latter to substitute for $W(x|y)$ in the former to obtain

$$rU(x) = \beta[\sigma(x|y) - (r + \delta)c] \tag{3.7}$$

where $\beta = \eta/(r + \delta + \eta)$. Now suppose y has been randomly selected to move first in the bargaining game. In response to the share left for her, x can walk out and continue searching, which carries the value $U(x)$, or she can reject this share and wait for another round of bargaining, which carries a value denoted $V(x)$. Note that the first mover y cannot hope to attain a better position than she currently has: at best, she will find herself as first mover again in a later meeting, be it with the same agent x or another agent of the same type. As delay is costly, y seeks to seize the opportunity and to ensure that x accepts her offer. In turn, x will accept any implicitly offered payoff $W^O(x|y)$ that satisfies

$$W^O(x|y) \geq \max[V(x), U(x)] \tag{3.8}$$

as she would otherwise reject the offer or walk out. When x moves first, y requires

$$W^O(y|x) \geq \max[V(y), U(y)] \tag{3.9}$$

When an offer is rejected, the same logic implies that the first mover in the second meeting seeks to ensure agreement, so that the second meeting can be expected to result in a match. By assumption 3.3, the second meeting happens at rate η, so that

$$rV(x) = \eta\,[W(x|y) - c - V(x)] \tag{3.10}$$

in the putative equilibrium. Solving equation (3.10) for $V(x)$ and equation (3.4) for $U(x)$ establishes that $V(x) = U(x)$. The bargaining game thus reduces to a variant of Rubinstein's (1982) set-up, and we have the following result:

Lemma 3.1 (Bargaining equilibrium). *Given truthful signals and given search strategies as in the putative equilibrium situation, the following strategies form the unique subgame-perfect equilibrium (SPE) of the bargaining game:*

(i) x always proposes

$$\pi^*(x|y) = \left(1 - \frac{\beta}{2}\right) f(x, y) + \beta(r + \delta)c \tag{3.11}$$

for herself; when y proposes $\pi(y|x)$ for herself, x always accepts if and only if $\pi(y|x) \leq \pi^*(y|x)$

(ii) y always proposes $\pi^*(y|x) = \pi^*(x|y)$ for herself; when x proposes $\pi(x|y)$ for herself, y always accepts if and only if $\pi(x|y) \leq \pi^*(x|y)$

Agreement is reached in the first round of bargaining.

Proof. Equations (3.8) and (3.9) uniquely determine equilibrium, as follows. Whenever y moves first, her optimisation problem is

$$\max \pi(y|x) \quad \text{s.t.} \quad W^O(x|y) \geq \max[V(x), U(x)]$$

With $V(x) = U(x)$, the constraint becomes $W^O(x|y) - U(x) \geq 0$. When match output is $f(x,y)$ and y takes a share $\pi(y|x)$ from the pot for herself, $f(x,y) - \pi(y|x)$ would be left for x. Therefore, in analogy to equation (3.5),

$$rW^O(x|y) = f(x,y) - \pi(y|x) - \delta[W^O(x|y) - U(x)]$$

Solving this for $W^O(x|y)$, we find that $W^O(x|y) - U(x) \geq 0$ if and only if

$$\begin{aligned} f(x,y) - \pi(y|x) - rU(x) &\geq 0 \\ \Leftrightarrow f(x,y) - \pi(y|x) &\geq \beta[\sigma(x|y) - (r+\delta)c] \end{aligned}$$

where the second line uses equation (3.7). Next substituting for $\sigma(x|y)$ from equation (3.6) and rearranging,

$$[f(x,y) - \pi(y|x)]\,\phi \geq \pi(x|y) - 2(r+\delta)c \tag{3.12}$$

where $\phi = (1 - \frac{\beta}{2})/\frac{\beta}{2}$. As y raises $\pi(y|x)$ the left-hand side of equation (3.12) linearly falls, while the right-hand side stays constant. Hence this constraint will hold with equality for the equilibrium value of $\pi(y|x)$. When x moves first, the constraint is analogously found as

$$[f(x,y) - \pi(x|y)]\,\phi \geq \pi(y|x) - 2(r+\delta)c \tag{3.13}$$

As binding constraints, equations (3.12) and (3.13) are two equations in two unknowns, so that they determine a unique equilibrium. To keep the procedure of the proof more generally applicable, we ignore the symmetry of this case. We thus rather solve equation (3.12) for $\pi(x|y)$ before substituting for $\pi(x|y)$ in equation (3.13).

Solving the resulting expression for $\pi(y|x)$, we obtain

$$\begin{aligned}
\pi(y|x) &= \frac{(1-\phi)\phi}{1-\phi^2} f(x,y) + \frac{1-\phi}{1-\phi^2} 2(r+\delta)c \\
&= \frac{\phi}{1+\phi} f(x,y) + \frac{1}{1+\phi} 2(r+\delta)c \\
&= \left(1 - \frac{\beta}{2}\right) f(x,y) + \beta(r+\delta)c
\end{aligned}$$

Then $\pi(x|y)$ can be found by substituting for $\pi(y|x)$ above and turns out to be the same as $\pi(y|x)$ in this symmetric case. Because both first-mover shares have been derived under the constraint that the second mover accepts, agreement is reached in the first round of bargaining. Finally, for a proof that this equilibrium is subgame-perfect, see Rubinstein (1982). □

The essence of the bargaining SPE is that each agent makes offers that leave the other indifferent, and each agent accepts offers that make her indifferent or better off: the first-mover takes a share $\pi^*(x|y)$ such that the second-mover share

$$f(x,y) - \pi^*(x|y) = \frac{\beta}{2} f(x,y) - \beta(r+\delta)c$$

is just enough to prevent the second mover from taking her outside option and from rejecting. The two indifference conditions, depending on who moves first, then together pin down a unique SPE. Finally, expected shares in the SPE are

$$\sigma(x|y) = \sigma(y|x) = \frac{1}{2}\pi^*(x|y) + \frac{1}{2}\left[f(x,y) - \pi^*(x|y)\right] = \frac{1}{2} f(x,y) \tag{3.14}$$

as one would expect when the bargaining game is symmetric.

3.4.2 Participation and steady state

To ensure that all agents engage in search, c must not be so high that $U(x)$ becomes negative for some x, since each agent can obtain 0 by refusing to search.

Lemma 3.2 (Participation). *Assumption 3.3 is necessary and sufficient for all* *agents to engage in search.*

Proof. As match output is the only source of utility in the model, agents who do not engage in search obtain payoff 0. Then agent x will only engage in search if $U(x) \geq 0$. By equation (3.7), this requires

$$c \leq \frac{\sigma(x|y)}{r+\delta} \quad \Leftrightarrow \quad 2c \leq \frac{f(x,y)}{r+\delta}$$

using equation (3.14). If this holds for $f(\underline{x}, \underline{y})$, as stated in assumption 3.3, then it will also hold for the output generated in any other match because $f(x, y)$ is increasing in x and y by assumption 3.1. $\square$

Hence all unmatched agents engage in search, so that $u(x) > 0$ and, by consequence, $l(x) > 0$ for all $x \in \Theta$. Combining equations (3.2) and (3.3), we obtain the pointwise steady state in the putative equilibrium:

$$\delta[l(x) - u(x)] = \eta u(x) \quad \forall x \in \Theta \tag{3.15}$$

3.4.3 Signals and beliefs

In this section, we examine whether any one agent has an incentive to unilaterally deviate from putative equilibrium by choosing another signal. Therefore, we take as given that all other agents signal truthfully, that all believe signals to be truthful, as well as the other components of the putative equilibrium. We proceed by identifying first the conditions under which each agent prefers her match in the putative equilibrium (henceforth the *equilibrium match*) to any other match available to her. From this, we infer under which conditions there will be no incentive to deviate from the truthful signal.

There are two reasons why we need to worry about false signals in the first place. First, because true types are only privately observable, agents can perfectly imitate agents of other types by sending their signal and then bargaining as they would.

Second, agents might just imitate another type's signal and then renege on it in the meeting. Since search frictions make switching to another meeting costly, the other agent in the meeting might still accept the match. For example, consider a rather high type y_H who matches with x_H in the putative equilibrium. If y_H finds she has been lured into a meeting with a type $x_L < x_H$ by a false signal, she will nevertheless grudgingly accept whenever her share of $f(x_L, y_H)$ is not so far below her expected share of $f(x_H, y_H)$ that the costs of another meeting would be justified. Therefore, there can in principle be an incentive to send false signals.

We first compare the equilibrium match to matches with lower types. Without loss of generality, let us take the perspective of some agent with a type $x_H > \underline{x}$, so that lower types necessarily exist. We thus want to compare being matched with $y_H = x_H$ to being matched with $y_L < x_H$. The expected present discounted values of these matches to x are $W(x_H|y_H)$ and $W(x_H|y_L)$, respectively. In the spirit of the one-deviation principle, x reverts to the strategies of the putative equilibrium after the deviation. Hence, the asset equations for both $rW(x_H|y_H)$ and $rW(x_H|y_L)$ in analogy to equation (3.5) depend on the same $U(x)$ and thus differ only in the expected shares. Solving these two asset equations respectively for $W(x_H|y_H)$ and $W(x_H|y_L)$, we therefore find that $W(x_H|y_H) > W(x_H|y_L)$ if and only if $\sigma(x_H|y_H) > \sigma(x_H|y_L)$, where $\sigma(x_H|y_H)$ and $\sigma(x_H|y_L)$ denote the expected share obtained by x_H in a match with y_H and y_L, respectively.

Thus suppose a type $x_H > \underline{x}$ signals to be of type x_L in order to meet a type y_L. Agent x_H may continue to behave like a type x_L so as to conform to the beliefs of y_L given that all other agents signal truthfully. Recall from section 3.4.1 that neither agent's signal implies a binding outside option. Hence the bargaining equilibrium described by lemma 3.1 will be reached in the first round of bargaining. Then the expected flow utility for x_H in the match with y_L is

$$\begin{aligned}\sigma(x_H|y_L) &= \frac{1}{2}\left[f(x_H, y_L) - \frac{\beta}{2}f(x_L, y_L) + \beta(r+\delta)c\right] \\ &\quad +\frac{1}{2}\left[f(x_H, y_L) - \left(1-\frac{\beta}{2}\right)f(x_L, y_L) - \beta(r+\delta)c\right] \\ &= f(x_H, y_L) - \frac{1}{2}f(x_L, y_L) \qquad (3.16)\end{aligned}$$

If x_H moves first, she leaves a share to y_L as if output was $f(x_L, y_L)$ and keeps the rest of the actual output $f(x_H, y_L)$. If y_L moves first, y_L takes the first-mover share of $f(x_L, y_L)$ for herself and x_H obtains the actual remainder. In an equilibrium match, by contrast, x_H would obtain

$$\begin{aligned}\sigma(x_H|y_H) &= \frac{1}{2}\left[\left(1-\frac{\beta}{2}\right)f(x_H, y_H) + \beta(r+\delta)c\right] + \frac{1}{2}\left[\frac{\beta}{2}f(x_H, y_H) - \beta(r+\delta)c\right] \\ &= \frac{1}{2}f(x_H, y_H) \qquad (3.17)\end{aligned}$$

Comparing $\sigma(x_H|y_L)$ and $\sigma(x_H|y_H)$, we find the following:

Lemma 3.3 (Matches with lower types). *In the putative equilibrium, strict supermodularity of $f(\cdot,\cdot)$ is necessary and sufficient for any agent $x \in \Theta$ to strictly prefer the equilibrium match to a match with a lower type in which she perfectly imitates the lower type.*

Proof. Any agent $x_H > \underline{x}$ will strictly prefer the equilibrium match to a match with a lower type y_L if $W(x_H|y_H) > W(x_H|y_L)$. As argued above, this is equivalent to

$$\begin{aligned}\sigma(x_H|y_H) &> \sigma(x_H|y_L) \\ \Leftrightarrow f(x_H, y_H) - f(x_H, y_L) &> f(x_H, y_L) - f(x_L, y_L) \qquad (3.18)\end{aligned}$$

using equations (3.16) and (3.17). Next, note that we can write

$$f(x_H, y_L) = f(y_H, x_L) \qquad (3.19)$$

because $x_H = y_H$ and $y_L = x_L$, and in turn

$$f(y_H, x_L) = f(x_L, y_H) \tag{3.20}$$

by symmetry of $f(\cdot,\cdot)$ (see assumption 3.1). Therefore of course

$$f(x_H, y_L) = f(x_L, y_H) \tag{3.21}$$

which we use to substitute out $f(x_H, y_L)$ on the right-hand side of equation (3.18) only, obtaining

$$f(x_H, y_H) - f(x_L, y_H) > f(x_H, y_L) - f(x_L, y_L)$$

By definition 3.3, strict supermodularity of $f(\cdot,\cdot)$ is necessary and sufficient for this to hold. Finally, the type $\underline{x}$ matches with $\underline{y}$ in the putative equilibrium, so that a lower type than in the equilibrium match does not exist in this case. $\square$

Now suppose that x_H signals to be of type x_L, thus meets a type y_L, but then wants to renege on the signal and somehow disclose her true type. We thus ask whether a deviation to a match with a lower type becomes worthwhile with this behaviour, and we find a negative answer:

Lemma 3.4 (Reneging in matches with lower types). *In the putative equilibrium, any agent $x \in \Theta$ would strictly prefer the equilibrium match to a match with a lower type if types were instantly observable in meetings.*

Proof. We want to prove that some $x_H > \underline{x}$ will strictly prefer the equilibrium match to a match with a type $y_L < x_H$ when the true x_H is observed before bargaining begins. In analogy to equations (3.8) and (3.9), x_H and y_L would respectively accept if

$$W^O(x_H|y_L) \geq \max[V(x_H), U(x_H)], \quad W^O(y_L|x_H) \geq \max[V(y_L), U(y_L)]$$

where $V(x_H)$ and $U(x_H)$ are determined by

$$rV(x_H) = \eta[W(x_H|y_L) - c - V(x_H)], \quad rU(x_H) = \eta[W(x_H|y_H) - c - U(x_H)] \quad (3.22)$$

and similarly for y_L. For comparison with agents' equilibrium matches, we recall their expected equilibrium shares as

$$\sigma(x_H|y_H) = \frac{1}{2}f(x_H, y_H), \quad \sigma(y_L|x_L) = \frac{1}{2}f(x_L, y_L)$$

by equation (3.14), and consequently from (3.7) that

$$rU(x_H) = \beta\left[\frac{1}{2}f(x_H, y_H) - (r+\delta)c\right], \quad rU(y_L) = \beta\left[\frac{1}{2}f(x_L, y_L) - (r+\delta)c\right] \quad (3.23)$$

First suppose that the outside option of x_H does not bind, $V(x_H) \geq U(x_H)$. With $U(x_H) > U(y_L)$ by equation (3.23), the outside option of y_L cannot bind when that of x_H does not bind, so that $V(y_L) \geq U(y_L)$. We then have the same situation as in lemma 3.1, except for labels. We thus know from lemma 3.1 that $\pi^*(x_H|y_L) = \pi^*(y_L|x_H)$ and therefore

$$\sigma(x_H|y_L) = \sigma(y_L|x_H) = \frac{1}{2}f(x_H, y_L)$$

When outside options do not bind, the bargaining game is thus still symmetric despite different types. Because $\sigma(x_H|y_H) > \sigma(x_H|y_L)$ we can conclude that x_H strictly prefers her equilibrium match when her outside option does not bind. Now suppose the outside option of x_H binds, $V(x_H) < U(x_H)$. Solving equation (3.22) respectively for $V(x_H)$ and $U(x_H)$, we obtain

$$V(x_H) = \frac{\eta[W(x_H|y_L) - c]}{r+\eta}, \quad U(x_H) = \frac{\eta[W(x_H|y_H) - c]}{r+\eta} \quad (3.24)$$

so that $V(x_H) < U(x_H)$ if and only if $W(x_H|y_L) < W(x_H|y_H)$. Hence x_H strictly prefers her equilibrium match also when her outside option binds. $\square$

Hence even if x_H could immediately convince y_L of her true type, x_H would strictly prefer the equilibrium match, as she does when she would have to imitate some lower type. Based on lemmas 3.3 and 3.4, we show below that higher types never have an incentive to deviate from the putative equilibrium to matches with lower types if $f(\cdot,\cdot)$ is supermodular, for any beliefs that lower types might hold about deviants. In turn, whenever a deviant causes bargaining to fail, the other agent thus knows that she faces a strictly lower type: for a weakly higher type, a deviation would be equilibrium-dominated. Next recall from assumption 3.2 that in expectation the same time $1/\eta$ elapses before another round of bargaining as before a meeting with a different agent. When bargaining fails due to a deviant, the other agent now prefers by lemma 3.2 to meet a different agent: she simply chooses her equilibrium match rather than a match with some strictly lower type after the same expected delay.[8] As this insight is central to our argument, we formally state and prove it:

Lemma 3.5 (Equilibrium-dominated strategies). *Let agents' beliefs place probability* 0 *on the occurrence of equilibrium-dominated actions, let $f(\cdot,\cdot)$ be strictly supermodular, and consider a meeting between some x and y in the putative equilibrium. If x deviates such that bargaining fails, y will correctly believe to face a lower type and will choose to walk away.*

Proof. We first establish that any agent $x_H > \underline{x}$ always prefers, for any beliefs of $y_L \leq x_H$, her equilibrium match to a deviation such that she meets a weakly lower type with whom bargaining fails. We have to consider all possible beliefs held by y_L about the potential match output $f(x,y)$ when bargaining fails:

(i) $g(y_L|h^y) = f(x_H, y_L)$ so that y_L believes to face the true type x_H. By lemma 3.4, x_H then strictly prefers her equilibrium match.

(ii) $g(y_L|h^y) > f(x_H, y_L)$ so that y_L overestimates potential match output. By the same argument as in the proof of lemma 3.4, y_L does not believe the outside

[8]This logic will also apply if a deviation is only detected after the start of the match: it can only be detected when agents' initial bargaining agreement breaks down, so that there is no basis for production when agents wait for the new round of bargaining required for renegotiation.

option of x to bind: if it did, x would have had to pursue an equilibrium-dominated strategy. Given that outside options do not bind, lemma 3.1 applies (which does not depend on symmetry): y_L optimally pursues a bargaining strategy with first-mover and second-mover shares respectively given by

$$\pi(y|x) = \left(1 - \frac{\beta}{2}\right) g(y_L|h^y) + \beta(r+\delta)c$$
$$f(x,y) - \pi(x|y) = \frac{\beta}{2} g(y_L|h^y) - \beta(r+\delta)c$$

Since $0 < \beta < 1$, both shares are unambiguously increasing in $g(y_L|h^y)$. As $g(y_L|h^y) > f(x_H, y_L)$, y_L demands higher shares than under (i). Because x_H strictly prefers her equilibrium match under (i), she still prefers her equilibrium match when y_L is more demanding.

(iii) $f(x_H, y_L) > g(y_L|h^y) > f(x_L, y_L)$ so that y_L underestimates potential match output but still believes to face a higher type. By lemma 3.3, x_H will strictly prefer her equilibrium match if her type is underestimated and believed to be x_L, as she then either has to imitate x_L or cause bargaining to fail again. By the same argument as under (ii), x_H then still prefers her equilibrium match when y_L believes to face a higher type than x_L and thus demands higher shares.

(iv) $g(y_L|h^y) = f(x_L, y_L)$ so that y_L believes to face the same type as her own type. If $x_H > y_L$, the argument under (iii) applies. If $x_H = y_L$, lemma 3.1 implies that x_H would have preferred reaching a bargaining agreement with y_L (which coincides with the equilibrium match).

(v) $g(y_L|h^y) < f(x_L, y_L)$ so that y_L believes to face a strictly lower type. As we consider a unilateral deviation from the putative equilibrium by x_H, y_L has sent a truthful signal, so that her type has been disclosed to x_H. By lemma 3.4, y_L then strictly prefers her equilibrium match to a match with x_H who is perceived as a lower type. Hence y_L walks away to meet another agent, and x_H does not gain from the deviation.

Hence the deviation in question is equilibrium-dominated for weakly higher types than y_L. Now requiring that agents' beliefs place probability 0 on the occurrence of equilibrium-dominated actions, y_L must believe to face a strictly lower type when bargaining fails, $g(y_L|h^y) < f(x_L, y_L)$. By the argument under (v), y_L thus walks away when bargaining fails. As we supposed that $y_L \leq x_H$, the entire reasoning applies to any $y \in \Theta$ including $y = \bar{y}$. $\square$

Let us now turn to the incentive for lower types to deviate to a match with a higher type. Without loss of generality, consider some agent with a type $x_L < \bar{x}$, so that higher types necessarily exist. Now we want to compare being matched with an exactly corresponding type $y_L = x_L$, as in the equilibrium match, to being matched with a higher type $y_H > x_L$. The lower type x_L has two possibilities: she can either perfectly imitate x_H, or she can signal to have type x_H in order to meet y_H but then renege on the signal. We have just shown that, if x_L reneges such that bargaining fails, y_H will walk away, so that x_L does not gain from the deviation. If x_L herself walks away, she does not gain from the deviation either. Yet x_L does not have any other way to renege on her signal. In particular, should x_L simply claim to have a lower type at some point during bargaining, y_H will have no reason to take such a claim seriously:

Lemma 3.6 (Irrelevant communication). *Any claims by agents to have a lower type than signalled are not credible.*

Proof. We have to show that an agent x who is not of lower type than y also has an incentive to make such claims. In particular, consider an agent of type $x = y$ with whom y matches in the putative equilibrium. From lemma 3.1, the first-mover and second-mover shares obtained by y in the equilibrium match are respectively

$$\pi^*(y|x) = \left(1 - \frac{\beta}{2}\right) f(x,y) + \beta(r+\delta)c, \quad f(x,y) - \pi^*(x|y) = \frac{\beta}{2} f(x,y) - \beta(r+\delta)c$$

As argued before, both shares are increasing in $f(x,y)$. By assumption 3.1, $f(\cdot,\cdot)$ is increasing in both arguments. Hence agent x has an incentive to downplay her type in order to make y propose and accept lower shares. $\square$

One can also show that any claims to have a higher type will not be taken seriously either if agents place probability 0 on the occurrence of equilibrium-dominated actions: by lemma 3.4, a type would thereby claim to have taken an equilibrium-dominated action. In sum, any agent $x_L < \bar{x}$ strictly prefers her equilibrium match to a deviation such that she meets a higher type and reneges on her signal. The only way that x_L can match with y_H is therefore to perfectly imitate x_H. Then the expected flow utility for x_L is

$$\begin{aligned}\sigma(x_L|y_H) &= \frac{1}{2}\left[f(x_L,y_H) - \frac{\beta}{2}f(x_H,y_H) + \beta(r+\delta)c\right] \\ &\quad + \frac{1}{2}\left[f(x_L,y_H) - \left(1-\frac{\beta}{2}\right)f(x_H,y_H) - \beta(r+\delta)c\right] \\ &= f(x_L,y_H) - \frac{1}{2}f(x_H,y_H) \end{aligned} \tag{3.25}$$

If x_L moves first, she has to leave y_H the second-mover share of $f(x_H, y_H)$ to avoid being found out and can thus take whatever is left of the actual output $f(x_L, y_H)$. If y_H moves first, y_H takes the first-mover share of $f(x_H, y_H)$ for herself and x_L obtains the actual remainder. By contrast, the expected flow utility for x_L from her equilibrium match would be

$$\sigma(x_L|y_L) = \frac{1}{2}f(x_L,y_L) \tag{3.26}$$

A comparison of $\sigma(x_L|y_H)$ and $\sigma(x_L|y_L)$ yields the following result:

Lemma 3.7 (Matches with higher types). *In the putative equilibrium, strict supermodularity of $f(\cdot,\cdot)$ is necessary and sufficient for any agent $x \in \Theta$ to strictly prefer the equilibrium match to a match with a higher type in which she perfectly imitates the higher type.*

Proof. Any agent $x_L < \bar{x}$ will strictly prefer the equilibrium match to a match with a higher type y_H if $W(x_L|y_L) > W(x_L|y_H)$, which is equivalent to

$$\begin{aligned} \sigma(x_L|y_L) &> \sigma(x_L|y_H) \\ \Leftrightarrow f(x_H, y_H) - f(x_L, y_H) &> f(x_L, y_H) - f(x_L, y_L) \end{aligned}$$

using equations (3.25) and (3.26). By equations (3.19) through (3.21), we can replace $f(x_L, y_H)$ on the right-hand side by $f(x_H, y_L)$. Strict supermodularity is necessary and sufficient for the resulting equation to hold. Finally, for the type $\bar{x}$, a higher type than in the equilibrium match does not exist. □

We have thus identified conditions under which each agent $x \in \Theta$ strictly prefers her equilibrium match to a deviation to any other match. Also recall that matching rates are the same across marketplaces, so that matching rates do not reverse this preference. Corollary 3.1 collects the implications of this section for agents' signals and beliefs:

Corollary 3.1 (Truthful signals). *Let agents' beliefs place probability* 0 *on the occurrence of equilibrium-dominated actions. Then strict supermodularity of $f(\cdot,\cdot)$ is necessary and sufficient for each agent $x \in \Theta$ in the putative equilibrium to strictly prefer a truthful signal $\tilde{x} = x$. The only beliefs consistent with truthful signals are $\psi(y|\tilde{y}) = l(y|\tilde{y}) = 1$ for all $y = \tilde{y}$.*

Proof. Choose and fix some arbitrary unmatched agent with a type $x \in \Theta$ and call this type x_E. To this exemplary type, a type y_E exactly corresponds. Given the strategies for meeting and bargaining in the putative equilibrium, an agent of type x_E matches with an agent of type y_E at rate η unless there is a deviation. Further, types x_E and y_E also meet unless there is a deviation. Hence, if x_E does not deviate, but sends a truthful signal $\tilde{x}_E = x_E$, it must be that $\tilde{x}_E \in R(y_E)$. Since $R(y_E) \equiv D(y_E) \cap \Omega(y_E)$, also $\tilde{x}_E \in D(y_E)$. Given that agents meet exclusively their own type in the putative equilibrium, $|D(y_E)| = 1$. Then we have $\tilde{x}'_E \notin D(y_E)$ and also $\tilde{x}'_E \notin R(y_E)$ for any non-truthful signal $\tilde{x}'_E \neq x_E$. In words, a deviating x_E cannot meet y_E, and therefore cannot match with y_E. Hence x_E has to signal truthfully to obtain her equilibrium match, which she strictly prefers to a deviation by the preceding lemmas. Because the agent of type x_E was arbitrarily chosen, the reasoning extends to any agent $x \in \Theta$. Finally, if all signals are truthful, then $l(y|\tilde{y}) = 1$ for all $y = \tilde{y}$, and agents' beliefs can only be consistent if $\psi(y|\tilde{y}) = 1$ for all $y = \tilde{y}$. $\square$

Hence each agent $x \in \Theta$ finds it optimal to signal truthfully because this is the only way to obtain her equilibrium match, which she most prefers. As all agents therefore indeed signal truthfully, only beliefs that signals are truthful can be consistent with equilibrium play.

In conclusion, this section has presented an extensive but essentially simple reasoning. We found that higher types will never deviate from the putative equilibrium to match with lower types if $f(\cdot,\cdot)$ is supermodular. An agent who detects a deviation should therefore believe to face a lower type. When she can choose between continued bargaining with a lower type and her equilibrium match, she prefers the latter. Lower types can then only match with higher types by imitating them, but they will not gain from such a deviation if $f(\cdot,\cdot)$ is supermodular. Then each agent prefers not to deviate and consequently finds it optimal to signal truthfully.

3.4.4 Meeting strategies

In this section, we take it as given that signals are truthful and concentrate on agents' optimal meeting strategies. Consider three types x_L, x_M, and x_H, with $\underline{x} \leq x_L < x_M < x_H \leq \bar{x}$. Suppose these types search in the same marketplace, so that each of them can meet with y_L, y_M, or y_H. We know from lemma 3.4 that each x_H would prefer a match with y_H to a match with y_M or y_L. Using truthful signals, the agents of type x_H can profitably set up a new marketplace where agents of type x_H exclusively meet each other. In the initial marketplace, they would also meet less preferred types, which is not offset by any advantage in meeting rates. By setting up an exclusive marketplace, the congestion externality imposed by less preferred types is avoided (see Jacquet and Tan (2007) for details of this logic). Intuitively, it is best to attend the marketplace where one meets only one's most preferred type and where no other types stand in the way.

Given that signals are truthful, the remaining types x_M and x_L can no longer meet with y_H, as they would have to send a false signal to join the marketplace where y_H can be met. Among the possible matches, x_M prefers by lemma 3.4 the match with y_M, so that all agents of type x_M now set up an exclusive marketplace, leaving the initial marketplace to the agents of type x_L. This logic applies to any marketplace with different types, so that all types have their own marketplace in equilibrium. We will formalise and generalise this logic in section to show that it does not only apply in the putative equilibrium, but always when signals are truthful. Formally, the optimal meeting strategy of an agent x is therefore

$$D(x) = \{\tilde{y} : \tilde{y} = x\}$$

Given the optimal bargaining strategies in lemma 3.1, every meeting in the putative equilibrium then leads to a match, as one would expect when truthful signals allow agents to know everything in advance.

3.4.5 Discussion

By way of summary, the preceding subsections have each shown a component of the putative equilibrium situation to hold, given the other components. We thus found the pointwise steady state in the PBE. Given a supermodular match production function and beliefs that rule out equilibrium-dominated actions, agents seek to meet only exactly corresponding types and match with them, so that the resulting equilibrium matching of agents is PPAM. All agents then also signal their types truthfully and correctly believe that all other agents signal truthfully. Our model thus leads to PPAM under the same weak condition as in Becker's (1973) frictionless model, despite two kinds of search frictions.

Let us clarify the role that supermodularity plays for our results. Since types are privately observable and nothing keeps agents from imitating other types, an agent may match incognito with any type she likes. However, because actual match output then differs from the match output suggested by the signals, the deviant will only remain incognito if she bears the necessary adjustment: she has to give up as much of her own share as is necessary to bridge the gap when actual output is lower (otherwise bargaining fails and the other agent walks out), and she quietly pockets the excess output when actual output is higher. To explain why a lower type x_L would then not match incognito with a higher type $y_H > x_L$, supermodularity is key: $f(x_H, y_H) - f(x_L, y_H)$ is the necessary adjustment when y_H otherwise matches with x_H in equilibrium, while $f(x_L, y_H) - f(x_L, y_L)$ is the extra output produced in comparison to the equilibrium match of x_L. With $f(x_L, y_H) = f(x_H, y_L)$ in the latter, as established by equations (3.19) through (3.21), the necessary adjustment will exceed the extra output if $f(\cdot, \cdot)$ is strictly supermodular. From the perspective of a lower type, any possible gains from higher output with a higher type are therefore more than outweighed by the costs from adjustment.

More generally, supermodularity has in effect assumed the role of a single-crossing property in our model. This way, we obtain a fully separating equilibrium even though signals are costless. Separation is therefore not driven by cost differences, but by differences in marginal productivity of the same agent over different matches. However, these differences in themselves do not deliver a separating equilibrium for a general type distribution and for unrestricted sets of signals that agents can choose from, as we found in previous versions of this chapter. Under the realistic assumption that true types are only privately observable, as in this model, supermodularity does deliver full separation for a general type distribution and unrestricted signal sets.

Sorting in the separating PBE is driven by a logic apparently new to the literature. In intuitive terms, agents are effectively bound by their signal, so that a low type can only choose between "being herself" in a match with an equally low type and behaving like a high type in a match with a high type. The most desirable option, "being herself" in a match with a high type, is not available. When behaving like a high type implies disproportionate sacrifices due to supermodularity, low types prefer matches with equally low types.

To put this into a real-world labour market context, suppose a low-skilled worker faces the choice between working at McDonalds and working at McKinsey. While McKinsey would presumably pay a significantly higher wage, the low-skilled worker would have to perform at McKinsey like her high-skilled colleagues. It is easy to imagine that the sheer effort and the extra hours needed to reach this performance outweigh the benefit of a higher salary, so that the low-skilled worker actually prefers working at McDonalds. Whenever this is the case, McKinsey does not even have to check whether applications are truthful, but would still meet only those who claim to be high-skilled. Indeed, this seemingly paradoxical behaviour that our model rationalises appears to be widespread practice in recruiting. It also makes sense in practice to dismiss applicants who are known to have lied in their application: as our model suggests, it appears easier to find a replacement rather than to disentangle lies from truth for such applicants, thereby determine their actual qualifications, and then adjust the job requirements to fit these qualifications. The next section discusses key properties of the separating equilibrium.

3.5 Equilibrium properties

3.5.1 Efficiency

The separating equilibrium we have identified is efficient in a number of important respects. First and foremost, search costs are minimised, both for each agent individually and overall: in equilibrium, truthful signals allow each agent to ensure that no meeting is wasted, but that every meeting she attends results in a match. Hence, whenever an agent x searches, she attends a meeting and also matches at the first opportunity, that is, after an expected search time of $1/\eta$. This is the absolute minimum because a meeting necessarily precedes a match. In a standard search model, each match would typically be preceded by a number of pointless meetings, and only by chance will the first meeting of an agent result in a match. Therefore, search costs in standard search models are at least as high from the individual perspective as in our model with truthful signals, and much higher in expectation as well as on aggregate. Second, note that all agents match in equilibrium so that there is no unrealised surplus left in the form of unmatched agents, on the contrary:

Corollary 3.2 (Output efficiency). *If the match production function is supermodular, PPAM will maximise aggregate output.*

A proof of this can be found in the appendix of Becker (1973). Standard search models, be it with or without supermodularity of the match production function, do generally not maximise aggregate match output, as they lead to a certain degree of mismatch instead of PPAM. Finally, not only is each agent matched at the first opportunity in the equilibrium we found, but each agent is also matched with the type she most prefers among the types in the matching set (as she only meets agents of this type). This again contrasts starkly with standard search models, where the type that an agent x expects in a match is the expectation over the set of types with whom a match is mutually acceptable, not the most preferred element of this set.

3.5.2 Stability

In this section, we examine whether the equilibrium matching we found is a *stable matching*. Associated with any matching is a vector $s = (\sigma_1, \ldots, \sigma_n)$ specifying for each matched agent i the expected flow utility σ_i obtained in this matching, where n is the total number of matched agents. (This is a slight abuse of notation because agents in our model are not countable.) The following definition is standard:

Definition 3.4 (Stable matching). *A matching is stable if the associated vector s of expected flow utilities satisfies $\sigma_i \geq 0$, $\forall i$ and there is no match between any agents x_i and y_j such that $\sigma(x_i|y_j) > \sigma_i$ and $\sigma(y_j|x_i) > \sigma_j$.*

While Becker (1973) used the concept of the *core* in his seminal work, we argue in the appendix that, if a standard definition of the core is adapted to our model, the set of matchings in the core and the set of stable matchings will coincide. A proof that PPAM is a stable matching would thus also prove that it is a matching in the core of our model. We find that supermodularity of the match production function is a sufficient condition here for PPAM to be a stable matching:

Corollary 3.3 (Stability of PPAM). *Given truthful signalling and individual strategies as in the putative equilibrium situation, strict supermodularity of $f(\cdot,\cdot)$ is necessary and sufficient for perfect positive assortative matching to be a stable matching.*

Proof. Suppose to the contrary that PPAM is not a stable matching. Then there must be a match between unequal types that is preferred by both types to matches with exactly corresponding types. However, by the proof of lemma 3.5, the higher type in any match between unequal types would always prefer a match with an exactly corresponding type when $f(\cdot,\cdot)$ is supermodular. Hence a match between unequal types that is preferred by both does not exist. Finally, lemma 3.1 implies together with assumption 3.1 that $\sigma_i \geq 0 \ \forall i$ in PPAM, so that no agent prefers to be single either. $\square$

A stable matching is a most unusual result in a model with search frictions. In standard search models, agents cannot search selectively and might thus be matched with any type from a certain range of types. Of course, many of these types are only accepted because search frictions make continued search undesirable. A stable matching cannot be expected to arise under such circumstances and is very unlikely to arise by chance whenever the number of different types is not trivially small. Stable matchings normally only arise in frictionless models. We attribute the reason that a stable matching is achieved here despite search frictions to the signals: they allow agents to pursue their search almost as if there were no search frictions.

Adachi (2003) shows for a fairly general search model that the set of equilibria will reduce to the set of stable matchings in a model à la Gale and Shapley (1962) if search frictions become negligible. Our result in this section qualifies this finding in so far as search frictions remain in our model because agents do not meet immediately ($\eta < \infty$) and incur costs from meetings ($c \geq 0$), and yet a stable matching results.

This suggests that frictions do not prevent a stable matching in a search model as long as they do not keep agents from meeting specifically chosen types.

3.5.3 Uniqueness

While we have shown that the putative equilibrium situation exists as a separating equilibrium, there might also be other separating equilibria. However, we find in this section that this separating equilibrium is unique. We know that signals are truthful in any separating equilibrium.[9] We have also found that agents in the putative equilibrium join marketplaces where they meet exclusively their own type. It turns out that agents *always* divide themselves into exclusive marketplaces when signals are truthful:

Lemma 3.8 (Market segmentation). *Agents will meet only their own type in any separating equilibrium.*

Proof. Suppose there is at least one marketplace in which, with truthful signals, agents do not only meet their own type, so that two or more types meet. Focus on the lowest productivity type y_L in this marketplace. This type must be the most preferred feasible type of some other type in the marketplace, otherwise the other types would exclude y_L from the marketplace to reduce congestion. However, y_L cannot be the most preferred type of itself, as all agents of this type would then set up their own exclusive marketplace to avoid congestion from other types. Hence y_L must be the most preferred type of some higher type x_H.

We will show that this cannot be a separating equilibrium. When x_H and y_L bargain over $f(x_H, y_L)$ (known to both following truthful signals), we have $V(x_H) \geq U(x_H)$ (otherwise, this would imply that x_H does not prefer y_L most). We do not determine $U(y_L)$; suffice to note that $W^O(y_L|x_H) \geq \max[V(y_L), U(y_L)]$ for y_L to accept when x_H moves first. As the outside option of x_H does not bind, $W^O(y_L|x_H)$

[9] We ignore separating equilibria where signals are not truthful yet still informative because they are linked by a one-to-one mapping to agents' true types, and this mapping forms the basis of agents' correct beliefs. Such equilibria would only be variants of equilibria with truthful signals and would thus not add anything to our argument.

is non-decreasing in $U(y_L)$. $V(y_L)$ is given when x_H and y_L bargain, so that the case $U(y_L) > V(y_L)$ requires a greater $W^O(y_L|x_H)$ (and thus a greater second-mover share for y_L) than when $U(y_L) \leq V(y_L)$. Similarly, the first-mover share for y_L is non-decreasing in her outside option $U(y_L)$. Hence, $U(y_L) \leq V(y_L)$ is the more favorable case for x_H. Focus on this case and note

$$rV(x_H) = \eta\left[W(x_H|y_L) - c - V(x_H)\right], \quad rV(y_L) = \eta\left[W(y_L|x_H) - c - V(y_L)\right]$$

which is the same as we found for $V(x)$ in the putative equilibrium above (see equation (3.10)), except for labels. Hence we can apply lemma 3.1, where outside options were not binding either, to the more favorable case for x_H here. Therefore, $\pi^*(x_H|y_L) = \pi^*(y_L|x_H)$ and consequently

$$\sigma(x_H|y_L) = \frac{1}{2}f(x_H, y_L) \tag{3.27}$$

in the more favorable case for x_H.

Now consider two agents x_H and y_H who match with y_L and x_L, respectively, and let us examine whether agents x_H and y_H both rather prefer the match with each other. First, suppose that the outside option of x_H does not bind, $V(x_H) \geq U(x_H)$, when x_H and y_H bargain over $f(x_H, y_H)$. Then also $V(y_H) \geq U(y_H)$ because $x_H = y_H$. This bargaining situation is therefore completely symmetric, implying $\pi^*(x_H|y_H) = \pi^*(y_H|x_H)$. Alternatively, one can write down the asset equations for $V(x_H)$ and $V(y_H)$ and note that lemma 3.1 again applies. We thus obtain $\sigma(x_H|y_H) = \frac{1}{2}f(x_H, y_H)$ when outside options are not binding. In this case, x_H strictly prefers the match with y_H to the match with y_L even in the more favorable case for x_H, as $\sigma(x_H|y_H) > \sigma(x_H|y_L)$ because $f(\cdot,\cdot)$ is increasing in its arguments by assumption 3.1. By the same reasoning, y_H also prefers the match with x_H to the match with x_L in this case.

Next, we have to show that the outside option of x_H never binds when x_H and y_H bargain. If it did, so that $V(x_H) < U(x_H)$, then this would imply $W(x_H|y_H) < W(x_H|y_L)$ as in the proof of lemma 3.4. We show this by contradiction: suppose it binds, $V(x_H) < U(x_H)$, then also $V(y_H) < U(y_H)$ and moreover $U(x_H) = U(y_H)$ because $x_H = y_H$. Provided there are gains from trade, this completely symmetric bargaining situation leads to $\sigma(x_H|y_H) = \frac{1}{2}f(x_H, y_H)$ (as can be verified by adapting the proof of lemma 3.1, maintaining $U(x_H)$ as an unknown). With this $\sigma(x_H|y_H)$ and $\sigma(x_H|y_L)$ from equation (3.27), we have

$$\begin{aligned} rW(x_H|y_H) &= \frac{1}{2}f(x_H, y_H) - \delta[W(x_H|y_H) - U(x_H)] \\ rW(x_H|y_L) &= \frac{1}{2}f(x_H, y_L) - \delta[W(x_H|y_L) - U(x_H)] \end{aligned}$$

in analogy to equation (3.5), so that $W(x_H|y_H) > W(x_H|y_L)$ because $f(x_H, y_H) > f(x_H, y_L)$. In analogy to equation (3.24) and noting that x_H meets y_L at best at rate η in a mixed marketplace, we can write

$$V(x_H) = \frac{\eta[W(x_H|y_H) - c]}{r + \eta}, \quad \sup U(x_H) = \frac{\eta[W(x_H|y_L) - c]}{r + \eta}$$

when x_H and y_H bargain. This establishes that $V(x_H) > \sup U(x_H)$, yielding the contradiction.

We have thus shown that x_H and y_H always prefer to match with each other than to match with y_L and x_L, respectively. To reduce congestion, y_L and x_L would therefore be excluded from the marketplace, which is possible because agents signal truthfully. Hence the proposed marketplace with two or more types cannot exist in a separating equilibrium. $\square$

Since agents only meet their own type in any separating equilibrium, they can only match with their own type. Therefore, PPAM is the unique matching that may result in a separating equilibrium of our model. We can now conclude more comprehensively:

Proposition 3.2 (Uniqueness). *Whenever it exists, the separating equilibrium described by the putative equilibrium is unique up to off-equilibrium beliefs.*

No formal proof is needed, as this follows from our earlier results. We know from lemma 3.8 that any separating equilibrium would have to lead to PPAM, so that other separating equilibria would have to differ in agents' signals, their beliefs, their strategies for meeting and bargaining, or in the steady state. However, there is only one way for each agent to signal truthfully. Then only one specification of beliefs about equilibrium actions will be compatible with the fact that all agents signal truthfully. When signals are truthful, lemma 3.8 means that the meeting strategy determined in section 3.4.4 is the unique optimal strategy. Section 3.4.1 identified the unique bargaining SPE in this context. Finally, the steady state conditions uniquely determine a density for the unmatched agents of each type. Hence, other separating equilibria can only differ in beliefs about off-equilibrium actions.

3.6 Conclusions

This chapter has introduced costless signals into search with transferable utility and frictions. We find a unique separating equilibrium characterised by perfect positive assortative matching, minimised search duration and search costs, and maximised overall match output. The non-random search in this separating equilibrium reflects the pervasive use of effective communication in real-world matching markets that facilitates search. Positive assortative matching in this equilibrium depends just on supermodularity of the match production function. Supermodularity simultaneously ensures sufficient complementarity and replaces a single-crossing condition. Our model thereby proposes a solution to the paradox in Shimer and Smith (2000): supermodularity as such is unambiguously a weaker condition than the conditions they identified. In fact, our particularly mild condition does not merely ensure positive assortative matching despite discounting and explicit search costs, but even perfect positive assortative matching.

Our results would therefore explain why sorting is much more frequent across many different real-world matching markets than one would expect, given the conditions for sorting that had been identified in previous theoretical models. This chapter thus suggests that allowing for more information in the search process enables search models to generate sorting more easily. We have found this for the realistic case that agents control these additional information flows and that they may manipulate them strategically. We would argue that sorting becomes more important as technological and societal progress favours specialisation. At the same time, many new means have appeared of effective and rapid communication that might, as in our chapter, support sorting. The combination of these two developments offers ample scope for further research.

3.A Equivalence of core and stable matching

In Becker (1973), a matching is an equilibrium matching whenever it is in the core. A standard definition of the core is offered by Telser (1978):

Definition (Core). *Call the set C of agents a coalition, and let $Z(C)$ give the highest possible sum of flow utility the coalition can obtain under the most adverse conditions, with $Z(\emptyset) = 0$. Associated with any matching is a vector $s = (\sigma_1, \ldots, \sigma_n)$. The matching is said to be in the core only if, for all legal coalitions C,*

$$\sum_{i \in C} \sigma_i \geq Z(C)$$

In words, the core is the set of matchings such that no legal coalition can ensure more flow utility for all its members than obtained in the matching. However, if only the sum of utility obtained by the coalition counts, the possibility of side payments within the coalition is implicitly assumed. Side payments are crucial in Becker's (1973) reasoning: an individual agent then always prefers, among all matches available to her, the match generating the highest match output, since her partner in this match will use the extra output to outbid any other potential match partners. Yet where output is divided through bargaining, an agent's share in the match generating the highest output may fall short of her share in another match.

We thus have to modify the definition of the core to ensure that each agent's σ_i weakly exceeds the utility she obtains in any legal coalition available to her. In our model, the only legal coalitions are those with $|C| \in \{0, 1, 2\}$. Then such a modified definition reduces to two requirements: σ_i has to weakly exceed agent i's utility of being single ($|C| = 1$), and no match ($|C| = 2$) is available to i in which she obtains strictly more than σ_i. When agents go to perfectly segmented marketplaces under truthful signalling, we can identify a match that is available to i with a match where the match partner is better off than in any other available match. Definition 3.4 makes the same requirements.

Chapter 4

The time trend in the matching function[1]

4.1 Introduction

The performance of labour markets has typically been described by matching functions as known from the work of Pissarides (see Pissarides (2000), for example). The standard matching function relates the number of matches H in a labour market, i.e. hirings, to the stocks of vacancies V and unemployed job seekers U: $H = m(V, U)$. It is almost always estimated using a Cobb-Douglas specification. In many cases, a time trend is also included to examine how labour market performance has changed over time, so that $H = m(V, U, t)$. Petrongolo and Pissarides (2001) report that a time trend was included in 7 of the 16 studies of the aggregate matching function that they survey. Overall, the empirical results in these studies clearly suggest that there is a highly significant negative time trend, implying that labour market performance appears to deteriorate over time. At the disaggregate level of occupations, Fahr and Sunde (2004) likewise find significant negative time trends in most cases, and not a single instance of a positive time trend. Indeed, deteriorating labour market perfor-

[1]A greatly improved version of this chapter is now available as IAB Discussion Paper 3/2012. It notably finds that the empirical matching function might simply reflect the law of motion for vacancies, such that one estimates this law of motion with omissions when the matching function is supposed to be estimated. The resulting omitted variable bias can exactly explain the estimated time trends in matching functions: the unbiased estimate is essentially zero and insignificant.

mance might have assumed the status of a stylised fact in labour economics, driven by findings of negative time trends and closely related findings on shifting Beveridge curves (see e.g. Jackman et al. (1989)).

The literature has not reached a consensus on why labour market performance appears to be steadily deteriorating. Probably, institutional changes play an important role (see Nickell et al. (2005) for an overview). For example, if the unemployment benefit becomes more generous, then unemployed job seekers can afford to reject job offers with comparatively low wages that they would otherwise have accepted. This way, a rise in the unemployment benefit leads to fewer matches than before at any given level of vacancies and unemployed job seekers. Empirically, this will register as a decrease in labour market performance. However, estimation procedures can account for institutional changes, and they seem to explain but a part of the negative time trends (see Petrongolo and Pissarides (2001)). At the same time, the forces that improve labour market performance may be harder to account for, although such forces most likely exist. For example, over the last two decades, the increasing use of the internet in job search and recruitment would tend to speed up matching at any given level of vacancies and unemployed. Therefore, the common finding of negative time trends has remained puzzling.

In this chapter, we examine whether the negative time trend arises as a merely statistical product: the estimate of the time trend may be biased downwards, so that a seemingly negative time trend appears when the actual time trend is zero or even positive. Such a negative time trend then would not correspond to any actual changes in labour market performance. We hypothesise specifically that estimated matching functions suffer from omitted variable bias, which would affect all estimated coefficients, including the time trend. To our knowledge, this line of inquiry has not been taken before. We consider three variables that are omitted by the standard matching function $H = m(V, U, t)$ although they might be relevant in the true matching process between unemployed and vacancies: other job seekers (employed or non-participating), current inflows into vacancies and unemployment, and the probability that encounters between unemployed and vacancies lead to matches. For each variable, we explore theoretically and empirically how omitting it from the matching function affects the estimated time trend. To this end, we build a comprehensive model of labour market flows, which allows us to construct important data series that are otherwise unavailable.

Our findings are directly related to the studies that have established the result of a negative time trend: the three variables we consider are all almost entirely absent from the aggregate matching functions listed in Petrongolo and Pissarides (2001). Omissions from the matching function as such have received some attention. Broersma and van Ours (1999) highlight the error that results when measures of matches and job seekers do not correspond, e.g. precisely because employed job seekers are omitted. Mumford and Smith (1999) argue that matches involving unemployed job seekers should not be analysed in isolation from matches involving employed job seekers. Sunde (2007) analyses the problems that arise when only registered vacancies are observed alongside unemployed job seekers. However, none of these papers explores the consequences for the time trend. Rather, where they at all propose any solutions to the problem of omitted variables, they rely on unique data sets and can therefore hardly be replicated.[2]

Perhaps closest to this chapter is work by Gregg and Petrongolo (2005). Their analysis includes measures of inflows into vacancies and unemployment (but does not include the other two variables we consider). Estimating unemployment and vacancy outflow equations, they do not find a negative time trend in the vacancy outflow equation. They then speculate that an increase over time of the vacancy stock relative to the vacancy inflow might be linked to the negative time trend in a standard matching function. Perhaps even more interestingly for us, the magnitude of the significant negative time trend they report for the unemployment outflow equation roughly halves as they move from the standard model of random matching to a model with inflows. Where we analyse inflows as a potential source of omitted variable bias in a matching function, we present a similar finding. We can also conclude that the omission of other job seekers does not seem to generate any part of the time trend,

[2]For example, to implement Sunde's (2007) approach, one needs to observe whether the job seeker in a match was previously unemployed or employed and whether the vacancy was registered.

as it remains unchanged when they are included. We likewise offer some indirect evidence that the omission of agents' selectivity is probably not responsible for any substantial part of the negative time trend.

Below we proceed as follows. In section 4.2, we describe our data, find a negative time trend using the common approach, and offer a critique of previously advanced explanations for this finding. In section 4.3, we propose an empirical model of labour market flows. We then use series constructed from the model to examine whether the time trend arises from the omission of job seekers beyond the unemployed. Section 4.4 similarly considers the omission of inflows into unemployment and vacancies as a potential source of omitted variable bias in the estimated time trend. Finally, section 4.5 explores whether there is evidence of bias from the omission of agents' selectivity (or "pickiness") before section 4.6 concludes.

4.2 The puzzle: negative empirical time trends

4.2.1 Data

We use two data sets. The first is the *Job Openings and Labor Turnover Survey* (JOLTS) from the U.S. Bureau of Labor Statistics.[3] Following the Bureau, JOLTS is a sample of around 16,000 observations collected on a monthly basis from establishments across the United States. These establishments may be firms in the private non-agricultural sectors or government bodies at the local, State, and Federal levels. Data collection started in December 2000 and continues to date. We will use the observations from January 2001 to February 2011, giving us $T = 122$ monthly observations. Here lie the first advantages of this data set for our purposes: we are not aware of any major institutional changes over this period that would likely reduce the efficiency of matching, thereby leading to a negative time trend. Rather, we can imagine that an ongoing shift towards internet-based recruitment and job search over this period increased labour market efficiency. Another advantage is the monthly frequency of the data, which should allow us to largely avoid the issue of aggregation

[3] The data used here were obtained after the revisions to JOLTS in early 2011.

bias and any consequences for the time trend.

Establishments participating in JOLTS report information on their hirings, vacancies, and separations. Following the definitions of the Bureau of Labor Statistics, hirings in period t (where the period is a month) are defined as the number of workers added to the firm's payroll in period t. This includes seasonal workers and rehired staff after a layoff at least a week earlier, but does not include staff from temporary help agencies or similar contractors. A vacancy in period t is defined as an unfilled position, part-time or full-time, on the last business day of period t. An unfilled position exists if a specific position is currently not held by anyone but could be taken up within 30 days, and if the firm engages in recruitment efforts to fill the position. This does not include positions that must be filled internally, nor positions for which somebody has been hired but work has not yet begun. Here lies a further advantage of the JOLTS data: these definitions and the census-like way of data collection ensure that measurements are good, even exceptionally good in the case of vacancies. Therefore, in contrast to many other data sets, measurement error is unlikely to be a major problem in JOLTS.

The second data set we use is the U.S. Current Population Survey (CPS), conducted by the Census Bureau and beginning in 1940. The information in the CPS comes from a monthly representative survey of around 60,000 U.S. households (thus around 110,000 individuals). Based on responses about activities of the household members during the reference week, which is normally the week including the 12th of the month, the individuals' labour market status according to CPS definitions is inferred. A person from age 16 who holds a job is classified as employed, be it full-time, part-time, or temporary work. A person from age 16 is classified as unemployed if the person does not currently hold a job but would be available for work and, unless on temporary lay-off, has actively sought work for at least four weeks. This definition matches the economic definition of unemployment rather well, so that measurement error should again not be a major issue. Finally, a person from age 16 neither classified as employed nor classified as unemployed is deemed not to be participating in the labour force. The sum total of individuals thus classified corresponds to the non-institutional civilian population from the age of 16, as people in institutions such as prisons and in the army are disregarded.

The CPS underwent a major re-design in 1994, but all the series we will use begin later and have been collected under the same standards. It is worth noting that the definitions used for these classifications appear to accord with the definitions used for the JOLTS data.[4] For example, workers on temporary lay-off are counted as unemployed in the CPS, and the recall of a temporarily laid-off worker is duly counted as a hiring in JOLTS. For both the JOLTS and the CPS data, table 4.1 offers descriptive statistics for the variables and derived variables that we will employ, over the period from January 2001 to February 2011. All levels in the table except H_t/U_t and V_t/U_t are given in thousands. From JOLTS, we use the series on hires (H_t), vacancies (V_t), and total separations (TS_t). E_t, U_t, and N_t denote CPS series on the stocks of workers who are respectively employed, unemployed, and not participating in month t. We obtain the series U_t^{QS} by adding CPS stocks of currently unemployed job leavers and job losers, and U_t^{LR} by adding currently unemployed labour force entrants and re-entrants. While we use several of the CPS series on the unemployment stock by unemployment duration, we report here only the most relevant, the stock $U_t^{<5}$ of currently unemployed individuals who have been in this state for less than 5 weeks.

Table 4.1: Descriptive statistics

Variable	mean	st. dev.	min.	max.
H_t	4829	870	2705	6243
V_t	3679	809	2038	5691
TS_t	4855	909	3059	8125
E_t	140669	3402	135701	146584
U_t	9262	2894	6023	15628
U_t^{QS}	5998	2228	3628	10866
U_t^{LR}	3265	692	2273	4966
$U_t^{<5}$	2787	231	2297	3522
N_t	77414	3932	70083	86216
H_t/U_t	0.579	0.209	0.178	0.972
V_t/U_t	0.449	0.187	0.134	0.945

[4]For further details and comparisons, data definitions for both JOLTS and CPS data can be obtained from http://www.bls.gov/bls/glossary.htm.

4.2.2 Some observations from a standard analysis

The most widely used matching function relates the flow of hirings H_t to stocks of vacancies V_t and unemployed job seekers U_t:

$$H_t = m(V_t, U_t)$$

where $m(\cdot)$ is the matching function. We focus on the canonical Cobb-Douglas specification of $m(\cdot)$. Then a matching function with a time trend may be written as

$$H_t = V_t^{\beta} U_t^{\alpha} e^{K+\gamma t+\epsilon_t} \tag{4.1}$$

where β is the elasticity of matches with respect to the stock of vacancies, α is that elasticity with respect to job seekers, K is a constant, γ is the coefficient of the time trend, and ϵ_t is random error. As the parameters and in particular K are time-invariant, only the time trend will pick up changes over time. When the constant is interpreted as the speed of matching (see e.g. Lindeboom et al. (1994)), the time trend γt will indicate how the performance of the labour market has changed over time. For example, a negative value for γ implies that over time fewer hirings result from given stocks of vacancies and unemployed, so that labour market performance deteriorates.

In logarithms, equation (4.1) returns a linear model equation:

$$h_t = K + \beta v_t + \alpha u_t + \gamma t + \epsilon_t \tag{4.2}$$

where the lower-case letters h_t, v_t, and u_t indicate logarithms. Given time series on H_t, V_t, and U_t, equation (4.2) can be estimated. As this equation is linear, OLS will give the best results provided the assumptions of the Gauss-Markov theorem hold. Line (1) in table 4.2 reports results for such a regression. As in all tables throughout the chapter, one star indicates a coefficient that is significant at the at 1% significance level, while two stars indicate significance at the 5% significance level. Because a strong element of seasonality may be expected in the data, we can also extend the model in equation (4.2) by a full set of monthly dummies to account for seasonality, with December as reference category:

$$h_t = K + \beta v_t + \alpha u_t + \gamma t + \sum_{i=month} \delta^i D^i + \epsilon_t \tag{4.3}$$

The results are reported in line (2) of table 4.2. While the results in line (1) suggest plausible values for the elasticities β and α, the results in line (2) do not make much sense: the negative coefficient for u_t would mean that a rise in job seekers leads to a fall in hirings, and also the coefficient for v_t is surprisingly low.

Table 4.2: Matching function regressions using OLS ($T = 122$)

dependent variable	v_t	u_t	θ_t	t	const.	seas. dum.	adj. R^2
(1) h_t	0.845*	0.302*	–	-0.0018*	-1.087	no	0.64
s.e.	0.106	0.106	–	0.0004	1.764		
(2) h_t	0.298*	-0.172*	–	-0.0008*	7.384*	yes	0.95
s.e.	0.059	0.055	–	0.0002	0.949		
(3) f_t	–	–	0.772*	-0.0016*	0.172*	no	0.93
s.e.	–	–	0.026	0.0004	0.023		
(4) f_t	–	–	0.750*	-0.0017*	-0.039	yes	0.98
s.e.	–	–	0.013	0.0002	0.021		

The results in line (1) support the hypothesis of constant returns to scale (CRS) in matching: doubling the inputs V_t and U_t in equation (4.1) will ceteris paribus also double H_t if $\beta + \alpha = 1$. Under this restriction, we can rewrite equation (4.1) as

$$\frac{H_t}{U_t} = \left(\frac{V_t}{U_t}\right)^\beta e^{K+\gamma t+\epsilon_t} \tag{4.4}$$

and then linearise this to obtain

$$f_t = K + \beta\theta_t + \gamma t + \epsilon_t \tag{4.5}$$

where $f_t = \ln(H_t/U_t)$ is the *job-finding rate* of the unemployed and $\theta_t = \ln(V_t/U_t)$ is the logarithm of *market tightness* V_t/U_t: the more vacancies there are relative to job seekers, the tighter is the labour market. Line (3) in table 4.2 gives the results for a regression based on this equation. They plausibly suggest an elasticity $\beta = 0.772$ for vacancies and $1 - \beta = 0.228$ for unemployed job seekers. This time, extending the model in equation (4.5) by monthly dummies does not lead to implausible results, as line (4) shows. Finally, and importantly for this chapter, all regressions that produced plausible estimates also indicate a highly significant negative time trend of about equal magnitude.

Table 4.3: Matching function regressions using IV ($T = 121$)

dependent variable	v_t	u_t	θ_t	t	const.	seas. dum.	R^2
(5) h_t	1.623**	0.981	–	-0.0032**	-13.537	no	0.52
s.e.	0.705	0.621	–	0.0013	11.327		
(6) h_t	0.203	-0.248	–	-0.0006**	8.826*	yes	0.95
s.e.	0.150	0.137	–	0.0003	2.434		
(7) h_t	0.884*	0.334*	–	-0.0019*	-1.686	no	0.66
s.e.	0.105	0.110	–	0.0004	1.788		
(8) h_t	0.361*	-0.110	–	-0.0009*	6.311*	yes	0.95
s.e.	0.060	0.058	–	0.0002	1.002		
(9) f_t	–	–	0.777*	-0.0016*	0.179*	no	0.93
s.e.	–	–	0.028	0.0003	0.025		
(10) f_t	–	–	0.754*	-0.0017*	-0.034	yes	0.99
s.e.	–	–	0.010	0.0002	0.025		

It is well known that the assumptions of the Gauss-Markov theorem might not apply to the model in equation (4.1), so that OLS should not be used. A lot of attention has been devoted to the issue that hirings are measured as the total flow in a period but then regressed on stocks that change during this period, not least as an immediate consequence of hirings (see Petrongolo and Pissarides (2001) for an overview). Another issue is endogeneity bias that might arise because firms choose to open more vacancies when labour market performance is particularly high, in the sense that ϵ_t is particularly large. Such endogenous firm behaviour would generate a correlation between ϵ_t and the explanatory variable V_t (see Borowczyk-Martins et al. (2011)).

In order to avoid such problems, one might use an instrumental variable (IV) approach. To this end, we employ a generalised method of moments estimator with a robust weight matrix (in all cases of IV estimation in this chapter). The regression in line (5) of table 4.3 estimates a model like in equation (4.2) but instruments v_t and u_t respectively by v_{t-1} and u_{t-1}, with which they are highly correlated. Yet the results are poor: the estimated elasticities appear too high, and that of u_t is not even significant at the 5% significance level. When monthly dummies are included in line (6), the estimated elasticities are low and insignificant, while that of u_t has even switched signs. The regression in line (7) instruments only u_t by u_{t-1} and produces results close to those in line (1) of table 4.2. When only v_t is instrumented, the results are equally implausible as in line (5). This could mean that u_{t-1} is a good instrument while v_{t-1} is not; but since the results in line (7) are close to those in line (1), it also appears that a good instrument only for u_t either does not help, or that there was no problem in regression (1) to begin with. In any case, also the results in line (7) give way to implausible results when monthly dummies are included in line (8).

Line (9) in table 4.3 estimates the specification with CRS in equation (4.5) but instruments θ_t by θ_{t-1}, and line (10) estimates this specification with monthly dummies. In both cases, the results are remarkably close to the results in lines (3) and (4) of table 4.2 based on OLS. The results also seem to change little with the inclusion of monthly dummies. The time trends are without exception negative and significant at the 1% significance level, with almost identical magnitudes. It thus appears that the CRS specification delivers plausible and particularly robust results across the estimations considered here. This chapter will therefore often focus on this specification.

In turn, the results for the matching function specifications that do not assume CRS are not satisfactory: only regressions not instrumenting for v_t and without monthly dummies have produced potentially acceptable results. Yet these regressions

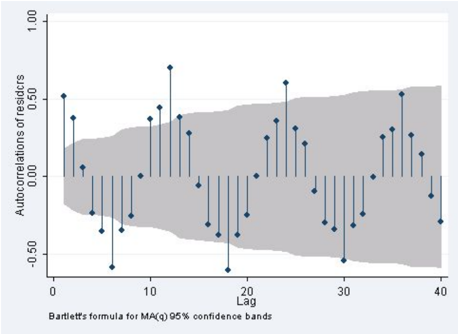

(a) Residuals of regression (3)

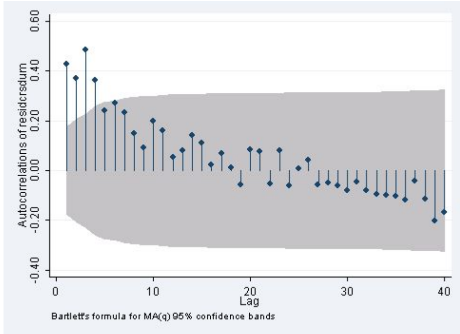

(b) Residuals of regression (4)

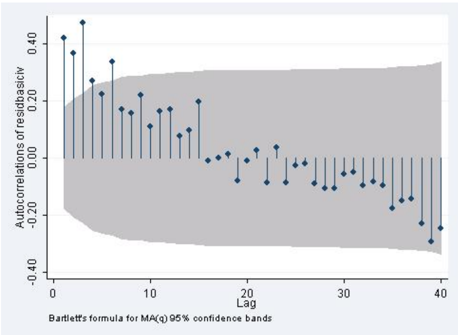

(c) Residuals of regression (6)

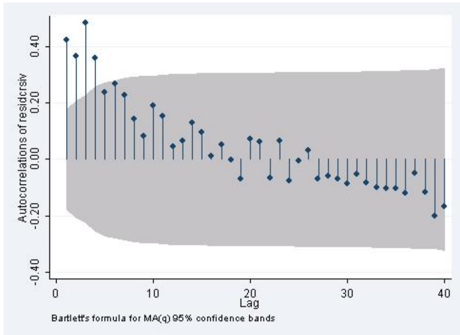

(d) Residuals of regression (10)

Figure 4.1: Autocorrelation of the regression residuals

are likely to suffer from the endogeneity bias mentioned above unless v_t is predetermined, which is not the case: the vacancy data in JOLTS counts job openings on the *last* business day of the month. This failure of the standard matching function to materialise in the JOLTS data should be in and of itself suspicious. The results for the CRS specification do not counterbalance this failure, as they might reflect as little as a positive correlation between the job finding rate and market tightness, rather than a structural relationship between hirings, jobseekers, and vacancies. Whatever problem affects the standard matching function here, it might also be the source of the negative time trend.

In particular, the residuals in all regressions considered above indicate that misspecification may be an issue. Panel (a) in figure 4.1 calls for the inclusion of monthly dummies to account for seasonality. Panel (b) shows that strong autocorrelation remains also with dummies included. While these two panels depict the residuals for the CRS specification, the residuals for the standard matching function behave very similarly. Such autocorrelation may reflect an omitted variable that affects hirings similarly over a number of periods. With IV regression, the residuals of the standard and CRS specifications are respectively depicted by panels (c) and (d). They suggest that the instrumentation did not fix the problem.

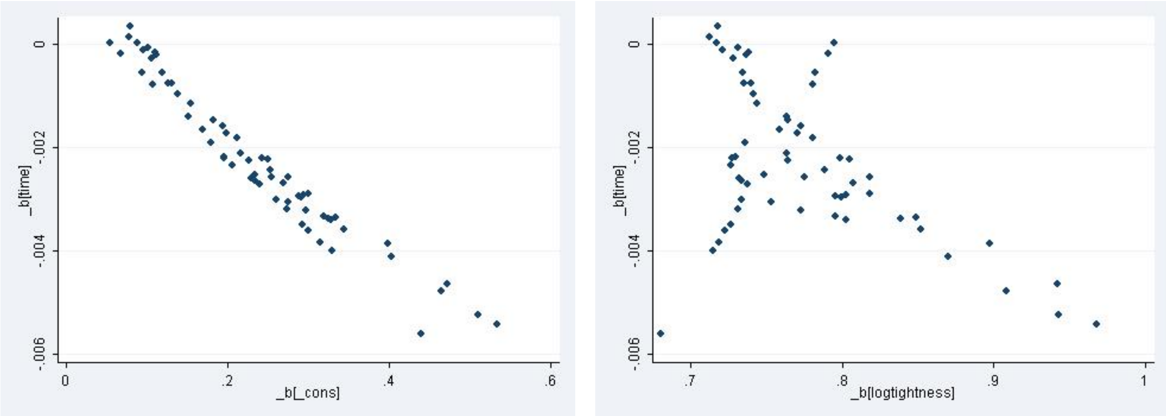

(a) $\hat{\gamma}$ plotted against the regression constant

(b) $\hat{\gamma}$ plotted against $\hat{\beta}$

Figure 4.2: Scatter plots from a rolling-window regression

To take a closer look, we perform the regression in line (3) as a rolling-window regression with 60 months covered by the window. Monthly dummies are disregarded here due to the reduced number of observations for each regression. Panel (a) in figure 4.2 shows a pronounced negative association between the estimated coefficient $\hat{\gamma}$ of the time trend and the estimated regression constant. We attribute this association to the mechanics of linear regression: given that the regression finds a negative time trend, its magnitude can be lower when a low constant sets a low starting level. It is more puzzling that panel (b) suggests a negative relationship also between $\hat{\gamma}$ and the coefficient $\hat{\beta}$ of market tightness. The well-known issues of simultaneity or endogeneity bias might not explain this association, as it is reinforced by the corresponding diagram for a rolling-window IV regression as in line (9). Instead, note that such a negative association is consistent with an omitted variable biasing $\hat{\gamma}$ and $\hat{\beta}$ respectively downwards and upwards. Finally, note the time trend appears to be negative in almost all instances. If taken at face value, this would indicate that labour market performance is almost always declining, not just its long-run average.

4.2.3 An assessment of previous hypotheses

It has often been suggested that a growing share of long-term unemployed is responsible for the negative time trend. Underlying this hypothesis is the idea that long-term unemployed may be less employable than short-term unemployed, due to the loss of skills over time or due to stigmatisation, or that long-term unemployed become discouraged and search less intensively than short-term unemployed. In any case, counting long-term unemployed towards job seekers with the same weight as short-term unemployed would overstate the measure of job seekers. If, in addition, the share of long-term unemployed grows over time, then job seekers are increasingly overstated over time. When there is no corresponding tendency in hirings, this generates a negative time trend in the standard matching function.

To account for the composition of unemployment, Blanchard and Diamond (1990) let only a proportion χ of long-term unemployed U_t^{LT} be perfect substitutes for short-term unemployed U_t^{ST}. The linearised standard matching function then becomes

$$h_t = K + \beta v_t + \alpha \ln(U_t^{ST} + \chi U_t^{LT}) + \gamma t + \epsilon_t$$

As it is not clear how χ should be determined, we take a radical approach here and set $\chi = 0$. We thus assume that an unemployed job seeker is a perfect substitute for a newly-unemployed up to a certain unemployment duration, and not a substitute at all from then on. We then consider three different cut-off durations: 5 weeks, 15 weeks, and 27 weeks. The relevant measures of unemployed job seekers are easily obtained from the CPS data on stocks of unemployed with duration of less than 5 weeks, 5 to 14 weeks, 15 weeks or longer, and 27 weeks or longer.[5] Even this radical approach, however, does not produce conclusive evidence. When the cut-offs 5 weeks, 15 weeks, and 27 weeks are used in OLS regressions without monthly dummies, the time trends are respectively −0.0009, −0.0011, and −0.0013 (all significant at the 2% significance

[5] With superscripts indicating the unemployment duration in weeks, we have

$$\begin{aligned} U_t^{<15} &= U_t^{<5} + U_t^{5-14} \\ U_t^{<27} &= U_t^{<15} + U_t^{\geq 15} - U_t^{\geq 27} \end{aligned}$$

level). First, it is worth noting that there is a significant negative time trend even when only newly unemployed are considered. The magnitude of this time trend does exhibit a tendency to increase as more long-term unemployed are counted in. Yet this result is not robust: as soon as monthly dummies are included, the magnitude of the time trend exhibits the reverse tendency, while the estimates of α are negative and signifcant, much like in line (2) of table 4.2. These latter results qualitatively persist when the vacancies and the respective unemployment measure are instrumented by their lagged values. Hence it appears that a growing share of long-term unemployed might at best explain a part of the time trend. This conclusion is in line with Blanchard and Diamond (1990) who find no evidence that short-term and long-term unemployed should at all be treated differently in the matching function.

Institutional changes are also often mentioned as a potential explanation for negative time trends. In particular, if unemployment insurance becomes more generous and job seekers become more selective as a consequence, then the same number of contacts in the labour market now leads to fewer matches. As the standard matching function does not distinguish between contacts and matches, this could generate a negative time trend. Yet attempts to account for the generosity of unemployment insurance, where quantifiable, have failed to explain the time trend (see Petrongolo and Pissarides (2001)). However, section 4.5 considers selectivity in more detail.

In a widely cited contribution, Bleakley and Fuhrer (1997) investigate the inward shift of the Beveridge curve for the U.S. in the late 1980s. Such shifts may capture the same changes as time trends in matching functions do. Indeed, among three potential explanations for the observed shift, Bleakley and Fuhrer (1997) first advance changes in matching efficiency. In this rare instance, their estimation of a standard Cobb-Douglas matching function produces a positive time trend, corresponding to the observed inward shift. While they conclude that shifts in the parameters K and α can equally well account for the time trend, they only speculate about the economic reasons behind this trend.

The second explanation they consider concerns changes in labour force growth: a permanently lower inflow of labour market entrants into unemployment would re-

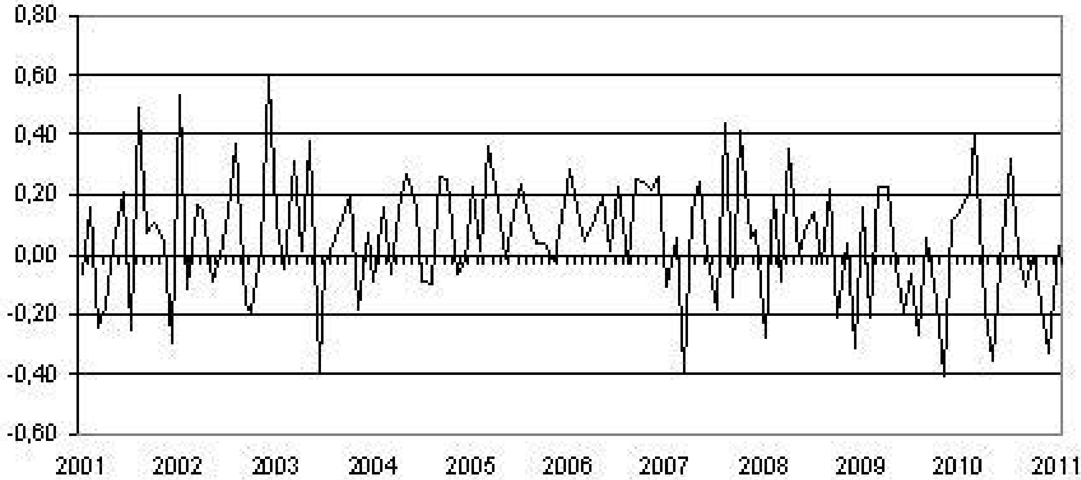

(a) Percentage change of the labour force, based on CPS data on employment and unemployment stocks

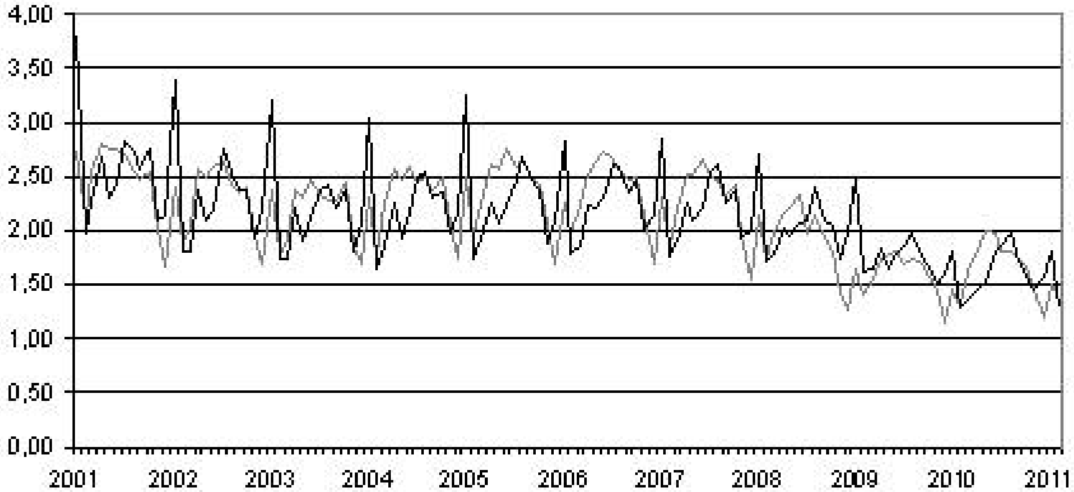

(b) Hirings (grey line) and separations (black line) in percent of the working-age population, based on flow data from JOLTS and a stock from the CPS

Figure 4.3: Graphical evaluation of the explanations in Bleakley and Fuhrer (1997)

duce unemployment levels at given vacancy levels and thus shift the Beveridge curve inwards. As Bleakley and Fuhrer (1997) offer a primarily graphical argument, we evaluate it using exactly corresponding graphs. Panel (a) in figure 4.3 suggests that this argument cannot explain the negative time trend we found above for the period January 2001 to February 2011. The slight tendency for labour force growth to fall over this period would, following this logic, again lead to a positive time trend.

Finally, Bleakley and Fuhrer (1997) point to a decrease in labour market "churning", measured as the magnitude of flows into and out of employment. If the magnitude of both flows decreases, there will be simultaneously lower levels of vacancies and unemployment, amounting to an inward shift of the Beveridge curve. Panel (b) in figure 4.3 shows that "churning" also decreased over the period considered here. For this to explain the negative time trend, the matching function would have to exhibit increasing returns to scale, so that a simultaneous fall in U_t and V_t disproportionately reduces H_t. The fragility of the standard matching function in our regressions above makes it hard to say much about the returns to scale. In fact, section 4.3.4 below finds evidence that, after an appropriate modification, an IV regression suggests constant returns to scale.

In conclusion, previous attempts to explain the phenomenon of negative time trends have had limited success. In this chapter, we theoretically and empirically examine three potential explanations that, to the best of our knowledge, have not been investigated before. We find that negative time trends can thus be explained in part and potentially in full.

4.3 Bias from the omission of other job seekers

4.3.1 Theory

While most studies take the stock of unemployed workers to be the stock of job seekers, we also seek to account for job seekers who currently hold a job or do not count towards the labour force. Because vacant jobs may be taken up just as well by such workers, the matching function will only be specified correctly if the stocks of employed and non-participating workers also enter. Several papers, for example Broersma and van Ours (1999) and Sunde (2007), have found that the estimated elasticities and returns to scale will be biased if only the stock of unemployed workers is used. However, the consequences for the estimated time trend have apparently not received any attention. To fill this gap, this section will investigate to what extent a negative time trend may result from omitted variable bias when employed and non-participating job seekers are omitted from the matching function regression.

We consider a more general matching function that relates the flow of hirings H_t to stocks of job seekers J_t and vacancies V_t:

$$H_t = m(J_t, V_t) = m(U_t + \phi_t(E_t + N_t), V_t)$$

The parameter ϕ_t denotes the average search intensity of employed and non-participating workers, relative to unemployed workers whose search intensity is normalised to 1. Assuming that $m(\cdot)$ again takes a Cobb-Douglas functional form, we have

$$H_t = V_t^{\beta} J_t^{\alpha} e^{K+\gamma t+\epsilon_t} = V_t^{\beta} \left[U_t + \phi_t(E_t + N_t)\right]^{\alpha} e^{K+\gamma t+\epsilon_t} \tag{4.6}$$

Since most studies in the literature either assume CRS or find evidence of CRS, let us also consider the special case $\beta + \alpha = 1$. Then dividing by U_t gives

$$\frac{H_t}{U_t} = \left[\frac{V_t}{U_t}\right]^{\beta} \left[\frac{U_t + \phi_t(E_t + N_t)}{U_t}\right]^{1-\beta} e^{K+\gamma t+\epsilon_t}$$

After linearising, this is

$$f_t = K + \gamma t + \beta\theta_t + (1-\beta)x_t + \epsilon_t \tag{4.7}$$

where

$$x_t = \ln \frac{U_t + \phi_t(E_t + N_t)}{U_t}$$

However, suppose a matching function is estimated that only includes unemployed workers and thereby omits x_t :

$$f_t = K' + \gamma' t + \beta' \theta_t + \epsilon_t \tag{4.8}$$

The textbook treatment of omitted variable bias in the context of multivariate regression leads to a simple expression for the expectation of the estimated coefficient $\hat{\gamma}'$:

$$\mathbb{E}(\hat{\gamma}') = \gamma + (1 - \beta)\tau \tag{4.9}$$

where τ is the coefficient of t in an OLS regression of x_t on all explanatory variables.

4.3.2 Measurement

In order to avoid this bias, stocks of employed and non-participating job seekers should also be included in the matching function. Since only some of these workers seek jobs, these stocks have to be weighted by ϕ_t, which requires information on flows between labour market states, in particular the various flows that generate matches. In principle, the CPS gross flow data provide such information for the U.S. labour market. These flows are obtained by comparing responses of individuals surveyed repeatedly in the CPS: if they report a different labour market status than the last time they were surveyed, this will count as a transition, and the sum of such transitions gives the respective flow. However, until recently the BLS did not even publish these flows because they were inconsistent with changes in the stocks.

As the gross flow data derive only from changes in labour market state, they cannot offer any series for the important flow of job-to-job transitions. In addition, two serious issues with the available flows have been discovered. Firstly, when individuals leave or enter the CPS sample, a comparison of their responses cannot be made, and these observations on transitions do not seem to be missing at random. Secondly, an individual's labour market status may be recorded incorrectly for some reason. While such classification errors might well cancel out in the data on stocks, they accumulate in the data on flows. For example, a constantly unemployed individual may be classified as employed once, generating a spurious transition, and is likely later classified correctly again, which generates a second spurious transition.[6]

Abowd and Zellner (1985) and Poterba and Summers (1986) find that the CPS gross flow data therefore very substantially overestimate the actual labour market flows, and they propose procedures for adjusting the data but admit that these procedures may themselves suffer from various problems. In any case, matching function regressions that rely on unadjusted CPS gross flow data should be treated very cautiously, including Bleakley and Fuhrer (1997) and studies as recent as Jolivet (2009). Nagypál (2008) points out that, due to the redesign of the CPS in 1994, the adjustment procedures can in fact not be used for data after 1994. She also finds that the incidence of classification error has even grown after the redesign.

Since the redesign the CPS has also asked employed individuals whether they still work for the same employer as before. Fallick and Fleischman (2004) use this information to construct a series for job-to-job transitions. From comparisons with other data sets, Nagypál (2008) finds that this method might overstate job-to-job transitions by as much as 30%. Moreover, the issue remains that missing observations are not missing at random. Finally, even if all these defaults were corrected, this will not allow us to construct series for the other labour market flows of interest. Unfortunately, weights for employed and non-participating job seekers in the matching function cannot be determined from job-to-job transitions alone.[7]

An alternative approach is to develop models that, given some more reliable data such as CPS stocks, allow to construct flow data. Shimer (2005b) builds a job ladder model that allows the construction of some flows (within employment and between unemployment and employment) from CPS data on stocks, but without using the CPS gross flow data. While the model is set in continuous time and thereby avoids time aggregation issues, its main drawback is the limitation to employment and unemployment as the only labour market states, which precludes in particular that matches may arise from non-participating job seekers. Adding non-participation would greatly

[6]These problems apparently extend to flow data based on similar surveys. For example, the Australian data used by Mumford and Smith (1999) suffer from much the same problems.

[7]Shimer (2005b) discusses a method that uses information from the March supplement to the CPS on jobs held over the previous year. It appears that this method involves still more problems.

increase this model's complexity and probably render it intractable.[8]

The approach we take in the next section is to build a simple but comprehensive accounting model of the labour market that treats all flows as unknowns. From the model's solutions, series for the unknown flows can be constructed using only the CPS data on stocks and some flow data from JOLTS. The JOLTS flow data are also collected by the CPS interviewers, but derive from changes in firms' payrolls and are thus free of the problems plaguing the CPS gross flow data. At the same time, the JOLTS data are equally regularly updated and equally easily available as CPS data. More generally, we believe that such an accounting model has the potential to give both a more comprehensive and a more accurate view of the labour market than a standard job ladder model. Indeed, Nagypál's (2005) results strongly suggest that standard job ladder models cannot explain the observed job-to-job transitions. We thus hope that our alternative approach stimulates further effort in this direction.

4.3.3 An empirical model of labour market flows

To explicitly account for job-to-job transitions and other flows, we assemble a comprehensive model of the empirical flows on the labour market. The model exists of accounting equations that link CPS data on stocks to JOLTS data on flows. In principle, this suffers from a time aggregation problem - if an agent's spell between two movements on the labour market is short enough, it will not have been recorded in the monthly CPS data. However, this issue has typically been discussed in the context of quarterly data that are much more affected by this problem.[9] As we would view, for example, sufficiently short spells in unemployment between two jobs more as an actual job-to-job transition, not too much seems lost when some spells shorter than a month are not counted. In any case, Burdett et al. (1994) find that effects from

[8]In Shimer (2007), a model that includes all three labour market states is built and solved numerically, but the model requires data on flows to construct instantaneous transition rates. That is, this model does not produce any series for the flows, it rather takes the CPS gross flow data as an input without any adjustment, despite the serious issues in these data. Similarly, a very limited model in Nagypál (2008) takes adjusted CPS gross flows as inputs to obtain estimates for two transition rates.

[9]Gregg and Petrongolo (2005) and Shimer (2005b) suggest methods to correct for time aggregation.

time aggregation are the lower the higher the frequency of the data. Their findings suggest that such bias is unlikely to be a problem for matching function regressions based on monthly data. Another potential problem could be that the CPS data do not count very short spells while the JOLTS series on flows do; we return to this issue below.

Each worker in the model over the age of 16 is always in one of three labour-market states: employment (E), unemployment (U), and non-participation (N). Let EU_t denote the flow of workers in period t from employment to unemployment, and similarly for all other flows between states. Recall that E_t, U_t, and N_t denote the CPS stocks of workers in each state in period t, and that we can distinguish within unemployment the stock U_t^{QS} of workers who are unemployed following quits or forced separations and the stock U_t^{LR} of unemployed labour-force entrants or re-entrants. As before, the length of period t is a month, in accordance with our monthly data. The JOLTS data in turn provide series on hires (H_t) and on three types of separations. Distinguishing between these types does not serve our purposes here, so that we will use only total separations (TS_t).

Those who are already in the labour force at age 16 or are still in the labour force beyond the age of 65 make up a small share of the total labour force at any point in time. The vast majority of workers in the labour force have come from state N after the age of 15 and return to that state before the age of 66. Hence, we can pretend that all flows into the working-age labour force originate in N, and all flows out of the working-age labour force lead to state N. Therefore, the growth of the working-age population directly only affects N and is then passed on to other states through the flows that originate in N. In principle, all unemployed must then have come from either employment or non-participation. However, U_t diverges slightly from the sum $U_t^{QS} + U_t^{LR}$ (by at most 2.2%). We thus adjust U_t^{QS} and U_t^{LR} such that their sum equals U_t, while the ratio U_t^{QS}/U_t^{LR} is preserved by the adjusted series U_t^{QS+} and U_t^{LR+}:

$$U_t^{QS+} = U_t^{QS}\left[1 + \frac{U_t - U_t^{QS} - U_t^{LR}}{U_t^{QS} + U_t^{LR}}\right], \quad U_t^{LR+} = U_t^{LR}\left[1 + \frac{U_t - U_t^{QS} - U_t^{LR}}{U_t^{QS} + U_t^{LR}}\right]$$

We also make the simplifying assumption that unemployed workers in U_t^{QS+} and U_t^{LR+} move into employment during period t at the same average rate UE_t/U_{t-1} (the job-finding rate), and also move into non-participation at the same average rate UN_t/U_{t-1}. We can then write down the following stock-flow equations:

$$\Delta U_t^{QS+} = EU_t - \frac{UE_t + UN_t}{U_{t-1}} U_{t-1}^{QS+} \tag{4.10}$$

$$\Delta U_t^{LR+} = NU_t - \frac{UE_t + UN_t}{U_{t-1}} U_{t-1}^{LR+} \tag{4.11}$$

Equation (4.10) says that the change $U_t^{QS+} - U_{t-1}^{QS+}$ equals the inflow of workers from quits or lay-offs during period t, EU_t, minus the outflow from the existing stock U_{t-1}^{QS+} at rate $(UE_t + UN_t)/U_{t-1}$. Equation (4.11) says the same for unemployed workers who were previously not participating.

We next link the JOLTS series to the flows in our model. As hirings include workers from any of the three labour market states, we know

$$H_t = EE_t + UE_t + NE_t \tag{4.12}$$

Likewise, following a separation, workers can move to any of the three labour market states:[10]

$$TS_t = EE_t + EU_t + EN_t \tag{4.13}$$

The difference $H_t - TS_t$ does not level with changes in E_t. This discrepancy cannot be due to time aggregation issues, as very short employment relations would contribute equally to H_t and TS_t and therefore leave $H_t - TS_t$ unaffected. Rather, H_t and TS_t reflect changes in firms' payrolls, while E_t is obtained from surveyed households and therefore also counts self-employed individuals.[11] Since self-employed do not match with firms, they should not feature in a matching model of the labour market. We therefore use $H_t - TS_t$ throughout to measure changes in employment, but we do not write down another stock-flow equation for these changes, as this equation would

[10]Given that changes in the population only affect state N in our model, deaths and emigration (taken into account by TS_t) are thought of as a part of EN_t.

[11]A CPS series for employment without self-employed is apparently not publicly available.

simply combine equations (4.12) and (4.13) and would thus not be independent.

To close the model, one can also write down a stock-flow equation for changes in non-participation. If we neglected for a moment that ΔN_t does not correspond to the flows in our model because N_t does not include the self-employed, we could write

$$\Delta N_t = \Delta P_t + (EN_t + UN_t) - (NE_t + NU_t)$$

where ΔP_t denotes the growth of the working-age population. However, this equation would not add an independent equation either. Since it should then hold that

$$\Delta U_t^{QS+} + \Delta U_t^{LR+} + H_t - TS_t + \Delta N_t = \Delta P_t$$

the equation for ΔN_t would be a combination of equations (4.10), (4.11), (4.12) and (4.13). Instead, we obtain another independent equation using the CPS series on the stock of workers who have been unemployed for less than five weeks, denoted $U_t^{<5}$. As our model does not count very short unemployment spells for the reasons mentioned above, this stock of newly unemployed represents the inflows into unemployment during the month:

$$U_t^{<5} = EU_t + NU_t \tag{4.14}$$

Equations (4.10) through (4.14) form a system of five equations in the following seven unknowns: EE_t, UE_t, NE_t, EN_t, EU_t, NU_t, and UN_t. To turn this into an exactly identified system, we make an identifying assumption. From the discussion in the last subsection, recall the CPS gross flow data that are plagued by response errors. Denote the CPS series for the flow from employment to non-participation by $\mathcal{EN}_t$ and the reverse flow by $\mathcal{NE}_t$. Further denote the respective proportions of these flows that are generated by response errors as p_t^{EN} and p_t^{NE}. Thus we have

$$\frac{\mathcal{EN}_t(1 - p_t^{EN})}{\mathcal{NE}_t(1 - p_t^{NE})} = \frac{EN_t}{NE_t}$$

Our identifying assumption is that $\mathcal{EN}_t$ and $\mathcal{NE}_t$ respectively overstate EN_t and NE_t by a roughly equal average proportion, $\mathbb{E}(p_t^{EN}) \approx \mathbb{E}(p_t^{NE})$. Then the ratio of the averages of $\mathcal{EN}_t$ and $\mathcal{NE}_t$ will tend to approximately the true ratio of the averages of EN_t and NE_t:

$$\operatorname*{plim}_{T\to\infty} \frac{\sum_{t=1}^{T} \mathcal{EN}_t}{\sum_{t=1}^{T} \mathcal{NE}_t} \approx \frac{\sum_{t=1}^{T} EN_t}{\sum_{t=1}^{T} NE_t} \tag{4.15}$$

Panel (a) in figure 4.4 shows how the ratio of $\mathcal{EN}_t$ to $\mathcal{NE}_t$ has evolved from 2001 to mid-2011. The ratio has fluctuated relatively little around a particularly stable average just above 1, and the ratio of the averages over this period is accordingly 1.05. Hence, one would expect the convergence in equation (4.15) to happen quickly, so that we can rely on the approximation already at comparatively small values of T. We therefore find that $\frac{1}{T}\sum_{t=1}^{T} EN_t \approx \frac{1}{T}\sum_{t=1}^{T} NE_t$. As nothing can be inferred about the variation of EN_t and NE_t, we use these averages in each period.

Under the identifying assumption, subtracting equation (4.13) from equation (4.12) leads the averages of EN_t and NE_t to cancel out:

$$H_t - TS_t = UE_t - EU_t \tag{4.16}$$

A single step has thus eliminated three unknowns. Together with equations (4.10), (4.11), and (4.14), we now have a system of four equations in four unknowns. In turn, it will be impossible to separately identify the three unknowns just eliminated, which is in fact not necessary: we will only need the sum $NE_t + EE_t$, which we denote by SE_t for short, to estimate a general matching function. This sum can be obtained from equation (4.12) once UE_t has been identified.[12]

To solve the system, we sum equations (4.16) and (4.14) and then use equation (4.11) to substitute for NU_t, which leads to

$$H_t - TS_t + U_t^{<5} = UE_t + \Delta U_t^{LR+} + \frac{UE_t + UN_t}{U_{t-1}} U_{t-1}^{LR+} \tag{4.17}$$

[12] Using a CPS series for the stock of discouraged workers (i.e. workers who are considered non-participating after a spell in unemployment that did not lead to a job), one could write down an additional independent equation that would allow to identify NE_t separately. However, we found that this approach produces most implausible results. One problem with this CPS series might be that it is unclear when those remaining in non-participation cease to be discouraged workers.

Next, we sum equations (4.16) and (4.10) to obtain

$$H_t - TS_t + \Delta U_t^{QS+} = UE_t - \frac{UE_t + UN_t}{U_{t-1}} U_{t-1}^{QS+} \tag{4.18}$$

We then solve equations (4.17) and (4.18) respectively for UN_t and equate the resulting expressions. Then one can solve for the only remaining unknown:

$$UE_t = H_t - TS_t + \frac{1}{1 + \frac{U_{t-1}^{LR+}}{U_{t-1}^{QS+}}} \left[U_t^{<5} - \Delta U_t^{LR+} + \frac{U_{t-1}^{LR+}}{U_{t-1}^{QS+}} \Delta U_t^{QS+} \right] \tag{4.19}$$

We next obtain EU_t from equation (4.16) as

$$EU_t = UE_t - H_t + TS_t = \frac{1}{1 + \frac{U_{t-1}^{LR+}}{U_{t-1}^{QS+}}} \left[U_t^{<5} - \Delta U_t^{LR+} + \frac{U_{t-1}^{LR+}}{U_{t-1}^{QS+}} \Delta U_t^{QS+} \right] \tag{4.20}$$

and subsequently NU_t from equation (4.14) as

$$NU_t = U_t^{<5} - EU_t = \frac{U_t^{<5}}{1 + \frac{U_{t-1}^{QS+}}{U_{t-1}^{LR+}}} + \frac{1}{1 + \frac{U_{t-1}^{LR+}}{U_{t-1}^{QS+}}} \left[\Delta U_t^{LR+} - \frac{U_{t-1}^{LR+}}{U_{t-1}^{QS+}} \Delta U_t^{QS+} \right] \tag{4.21}$$

From equation (4.10) (or equivalently from equation (4.11)), we find UN_t:

$$\begin{aligned} UN_t \quad &= \quad \frac{U_{t-1}}{U_{t-1}^{QS+}} \left[EU_t - \Delta U_t^{QS+} \right] - UE_t \quad = \quad TS_t - H_t \\ &+ \quad \frac{1}{U_{t-1}^{QS+} + U_{t-1}^{LR+}} \left[(U_{t-1} - U_{t-1}^{QS+})(U_t^{<5} - \Delta U_t^{LR+}) - (U_{t-1} + U_{t-1}^{LR+}) \Delta U_t^{QS+} \right] \end{aligned}$$

Finally, given the results for UE_t and EU_t, equations (4.12) and (4.13) respectively return

$$EE_t + NE_t = H_t - UE_t, \quad EE_t + EN_t = TS_t - EU_t \tag{4.22}$$

The time series for the flows that we can thus construct are reported in full in the appendix, while table 4.4 gives their descriptive statistics (all levels in thousands). Panel (b) in figure 4.4 depicts the series for UE_t, EU_t, and SE_t. It is worth noting that the correlation between SE_t and market tightness is 0.73, in line with theoretical results on procyclical endogenous search by employed job seekers (see Burgess (1993), for example).

We can now investigate how it matters that the JOLTS series for H_t and TS_t also count very short spells while the CPS series might not: the analytic solutions for UE_t and EU_t include the term $H_t - TS_t$, so that such spells should cancel out, as every transitory spell in employment counts equally towards H_t and TS_t. Only the values for $EE_t + NE_t$ and $EE_t + EN_t$ are therefore affected and might thus be overestimated relative to the other flows.

4.3.4 Empirical evaluation

The series constructed from our empirical model enable us to estimate matching functions that account for all job seekers. We thereby offer a new empirical strategy to make such an estimation possible for the U.S. labour market. By using data on CPS stocks and JOLTS flows, this strategy is not limited to a rare data set, as in Anderson and Burgess (2000), nor does it, as in Jolivet (2009), rely on the CPS gross flow data (unadjusted for the problems discussed above). Like these contributions, however, we assume that the relative rates at which different job seekers transit into employment depends on their relative search intensities. Normalising the search intensity of unemployed job seekers to 1, the average search intensity ϕ of employed and non-participating job seekers in period t can then be determined from our model:

$$\frac{UE_t}{U_{t-1}} = \frac{EE_t + NE_t}{\phi_t(E_{t-1} + N_{t-1})} \quad \Leftrightarrow \quad \phi_t = \frac{EE_t + NE_t}{UE_t} \frac{U_{t-1}}{E_{t-1} + N_{t-1}} \tag{4.23}$$

Having calculated ϕ_t, we can now assemble the measure J_t of total job seekers and estimate the log-linear model in equation (4.6). Line (11) in table 4.5 reports the OLS results.[13] The insignificance of the coefficient for j_t is counter-intuitive. Plausible

[13] Since ϕ_t is calculated using lagged values of E_t and N_t, the timing is only aligned in the IV regressions in table 4.5. However, aligning the timing in the OLS regressions produces only marginally different results. Further, while most of the following regressions include constructed data, we do not adjust the standard errors, as any such attempt would inevitably increase them. Then estimated time trends might be insignificant only due to the higher standard errors, while they might be significant if collected data were used. Finally, note that all regressions are now based on fewer observations because recurrent outliers in the JOLTS data (December and January) generate extreme outliers in UE_t and thus in ϕ_t. Using an interpolated series for UE_t led results to

Table 4.4: Descriptive statistics for the constructed series

Variable	mean	st. dev.	min.	max.
U_t^{QS+}	5997	2229	3615	11038
U_t^{LR+}	3265	692	2260	4937
UE_t	1779	695	-110	3182
EU_t	1786	251	1235	2668
NU_t	1002	115	691	1148
UN_t	946	746	-332	3205
SE_t	3042	945	1036	5680

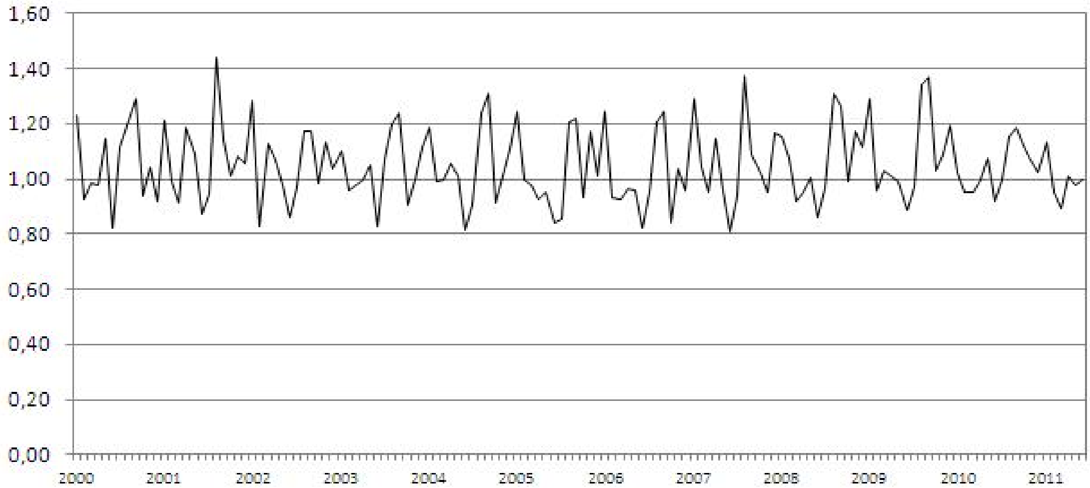

(a) The ratio of CPS gross flows from employment to non-participation ($\mathcal{EN}_t$) and from non-participation to employment ($\mathcal{NE}_t$)

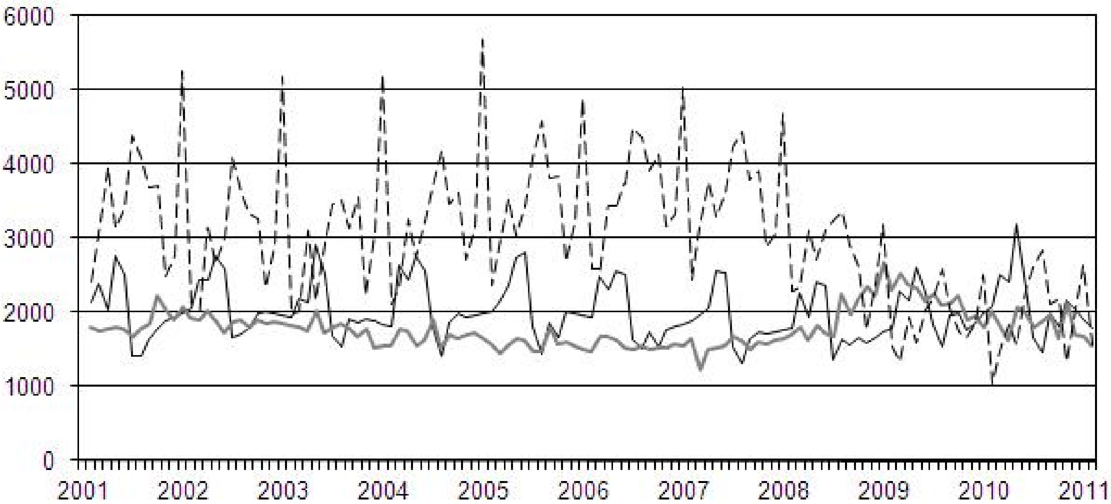

(b) The constructed series: the flows from non-participation and employment to employment (SE_t; top black line), from unemployment to employment (UE_t with interpolation of low outliers; bottom black line), and from employment to unemployment (EU_t; grey line). Levels in thousands

Figure 4.4: Assumption and result of the empirical model of labour market flows

significant coefficients are estimated instead in line (12) where v_{t-1} and j_{t-1} are respectively used as instruments for v_t and j_t.[14] These results support the hypothesis of CRS. Therefore, we can impose the restriction $\beta + \alpha = 1$ and estimate the model

$$\ln \frac{H_t}{J_t} = K + \gamma t + \beta \ln \frac{V_t}{J_t} + \epsilon_t$$

which is done in line (13) of table 4.5 using OLS. In line (14), the same model is estimated using $\ln(V_{t-1}/J_{t-1})$ as an instrument for $\ln(V_t/J_t)$. In both regressions, the estimates of β are significant and plausible. Yet the plausible estimation results in lines (12) through (14) also include exactly the same estimate for γ as we found using only unemployed job seekers (see section 4.2.2). Hence these results suggest that the omission of employed and non-participating job seekers does not generate the time trend. This conclusion is confirmed when we estimate equation (4.7) with seasonal dummies and still obtain $\hat{\gamma} = -0.0017$. Using the result $1 - \hat{\beta} = 0.2647$ from this regression to quantify the bias, equation (4.9) gives the expectation of the estimated time trend in the specification without x_t as -0.0015, so that there is practically no bias from the omission of other job seekers.

Table 4.5: Matching function regressions with all job seekers

estimator and T	dependent variable	v_t	j_t	$\ln(V_t/J_t)$	t	const.	seas. dum.	R^2
(11) OLS	h_t	0.498*	0.060		-0.0011*	3.801*	yes	0.93
101		0.034	0.061		0.0002	0.810		
(12) IV	h_t	0.612*	0.324*		-0.0016*	0.295	yes	0.90
90		0.050	0.095		0.0002	1.321		
(13) OLS	$\ln(H_t/J_t)$			0.647*	-0.0017*	-0.294*	yes	0.97
101				0.018	0.0002	0.037		
(14) IV	$\ln(H_t/J_t)$			0.634*	-0.0017*	-0.292*	yes	0.97
90				0.018	0.0002	0.037		

deteriorate, so that these outliers were simply dropped.

[14] To replicate these exact results, insignificant dummies for July, October, and November should be dropped.

4.4 Bias from the omission of flows

4.4.1 Theory

A growing number of papers suggest that the flows into job seekers and vacancies should also feature in the matching function, not just the stocks. The reason for the coexistence of vacancies and unmatched job seekers might be that no mutually acceptable matches can be formed among these vacancies and job seekers. If this is the case, then the inflows of new vacancies and job seekers are central to the matching process: existing job seekers match with the flow of new vacancies, while existing vacancies match with the flow of new job seekers. Where such stock-flow matching happens, a canonical matching function as in equation (4.1) is likely to suffer from omitted variable bias because it neglects the flows of new vacancies and job seekers, and this might affect the estimated time trend.

From their model of stock-flow matching, Coles and Smith (1998) obtain a log-linear matching function that includes stocks as well as inflows. Their analysis suggests the model

$$h_t = K + \beta v_t + \alpha u_t + \gamma t + \eta \tilde{v}_t + \zeta \tilde{u}_t + \epsilon_t \tag{4.24}$$

where $\tilde{v}_t$ and $\tilde{u}_t$ respectively denote the logarithm of the inflow into vacancies and unemployed job seekers. If $\eta \neq 0$ or $\zeta \neq 0$, the estimate for γ when flows are omitted will be biased and inconsistent:

$$\mathbb{E}(\hat{\gamma}') = \gamma + \eta \tau_{\tilde{v}} + \zeta \tau_{\tilde{u}} \tag{4.25}$$

where $\tau_{\tilde{v}}$ and $\tau_{\tilde{u}}$ are the coefficients of t in regressions of $\tilde{v}_t$ and $\tilde{u}_t$, respectively, on all included explanatory variables.

To the best of our knowledge, only Gregg and Petrongolo (2005) discuss time trends in the context of stock-flow matching. Their somewhat different approach is to estimate separate model equations for the outflows from the stocks of vacancies and unemployment. They argue that stock-flow matching implies the following for the flow from unemployment into jobs:

$$UE_t = g(U_t, V_t, \tilde{V}_t, t) \tag{4.26}$$

for some function $g(\cdot)$ that is increasing in U_t, V_t, and $\tilde{V}_t$. As explained in Coles and Petrongolo (2008), $\tilde{U}_t$ is not included because newly unemployed job seekers are expected to match primarily with existing vacancies, thus hardly affecting job seekers in the stock of unemployment who match primarily with the flow of vacancies. We can estimate equation (4.26) using the series for UE_t from our model, while Gregg and Petrongolo (2005) can apparently not distinguish between UE_t and UN_t in their data and therefore have to use the sum as dependent variable.

We can also extend the stock-flow reasoning to non-participating and employed job seekers here, using our constructed series. Let $O_t = \phi_t(E_t+N_t)$ denote these other job seekers. Those in O_t who are at risk of becoming unemployed can be thought of as the inflow into unemployment; they match primarily with existing vacancies in V_t, typically before they are counted towards the stock of unemployed. To the extent that they compete with unemployed job seekers for vacancies in V_t, U_t might also play a role. Others in O_t remain in their current status and wait for a suitable vacancy to appear; they thus match primarily with vacancies in $\tilde{V}_t$. The outflow from O_t, which is SE_t, should obey

$$SE_t = q(O_t, U_t, V_t, \tilde{V}_t, t) \tag{4.27}$$

for some function $q(\cdot)$ that is non-increasing in U_t but increasing in O_t, V_t, and $\tilde{V}_t$.

Gregg and Petrongolo (2005) find significant negative time trends only in the outflow from unemployment, and the magnitude appears to halve as they switch from a model of random matching to a model of stock-flow matching. They attribute the negative time trend in standard matching functions to a rise over time of V_t relative to $\tilde{V}_t$. This view thus seems analogous to the familiar argument on the share of long-term unemployed (see section 4.2.3): V_t largely consists of vacancies that have not been taken up before and are thus less likely to be taken up than those in $\tilde{V}_t$. A growing share of 'long-term vacancies' would bias the time trend downwards in a canonical matching function that, by omitting $\tilde{V}_t$, does not distinguish between V_t and $\tilde{V}_t$.

4.4.2 Empirical evaluation

We want to estimate equations (4.24) through (4.27) for the U.S. labour market. The empirical model we built in section 4.3.3 allows us to do this without additional data. Recall from equation (4.14) that we identified the total inflow into unemployment, $\tilde{U}_t$, with $U_t^{<5}$. Next, the spirit of the model implies a further accounting equation:

$$\Delta V_t = \tilde{V}_t - H_t$$

Assuming that vacancies are only closed when filled, the change in vacancies is given by the difference in hires and the inflow into vacancies. As V_t and H_t are known from the JOLTS data, this implies a series for $\tilde{V}_t$. We can now estimate equation (4.24), and the results are reported in line (15) in table 4.6. While the estimated coefficient of $\tilde{v}_t$ is positive and significant at the 1% significance level, that of v_t is insignificant, and those of u_t and $\tilde{u}_t$ are even negative and significant at the 1% significance level. This finding for $\tilde{u}_t$ appears recurrent in estimations of stock-flow matching functions, and Coles and Smith (1998) interpret it as a crowding-out effect among the unemployed.[15]

The estimated coefficients of stocks and flows in line (15) might in principle add up to one, which would justify a CRS specification. Although this is not the case here, it might still be worthwhile including the CRS specification because it has been considered in the literature (see for example model 2 in Gregg and Petrongolo (2005)). We therefore estimate the model

$$f_t = K + \beta\theta_t + \gamma t + \eta \ln \frac{\tilde{V}_t}{U_t} + \zeta \ln \frac{\tilde{U}_t}{U_t} + \epsilon_t$$

[15]To replicate their exact specification, one can repeat the regression in line (15) with the logarithm of the re-employment probability $\ln(H_t/U_t)$ as dependent variable. Then the estimated coefficient of u_t is strongly negative and significant while all other estimates remain unchanged, thus giving qualitatively similar results as in Coles and Smith (1998).

in line (16) of table 4.6. The regression produces estimates of β, η, and ζ that are positive and significant at the 1% significance level.

Next, we allow for both stock-flow matching and other job seekers. In line (17), we estimate the model in equation (4.24) with u_t replaced by j_t. The results appear more plausible than in line (15): the estimated coefficients of v_t and $\tilde{v}_t$ are positive and significant, while those of j_t and $\tilde{u}_t$ are negative and significant.[16] For the case that the coefficients of stocks and flows add up to one here, we again consider a CRS version:

$$\ln \frac{H_t}{J_t} = K + \beta \ln \frac{V_t}{J_t} + \gamma t + \eta \ln \frac{\tilde{V}_t}{J_t} + \zeta \ln \frac{\tilde{U}_t}{J_t} + \epsilon_t$$

The results, reported in line (18), are qualitatively the same as in line (16).

While the estimated time trend in line (15) is very low and only significant at the 5% significance level, lines (16) to (18) report negative estimated time trends that are again significant at the 1% significance level. Their magnitude is about half of the magnitude we found using either a standard approach or a model with all job seekers. This indicates that a substantial part of the estimated time trend in standard matching functions may be due to the omission of inflow measures. We reach the same conclusion based on equation (4.25): taking the estimate -0.00028 in line (15) as the true value of γ, we obtain $\mathbb{E}(\hat{\gamma}') = -0.00078$ for the time trend in a standard matching function. That is, if inflow measures are omitted, the resulting bias will increase a negative time trend of very small magnitude in the stock-flow matching function to roughly half the magnitude estimated for the standard matching function. However, the other half of the time trend in the standard matching function remains as unexplained as the fact that there is apparently a negative time trend also in the stock-flow matching function.

In lines (19) and (20), we estimate the models in equation (4.26) and (4.27), respectively. For simplicity, we assume that $g(\cdot)$ and $q(\cdot)$ can be written in a reduced form that is also log-linear. With the exception of V_t, all variables that were expected to play a role have significant estimated coefficients, also signed as expected. The

[16] We do not consider the inflow into J_t here because those who have just started a new job are perhaps least likely to look for another job.

Table 4.6: Stock-flow matching function regressions (OLS with seasonal dummies)

dep. var. and T	v_t	u_t	θ_t	j_t	$\ln \frac{V_t}{J_t}$	O_t	$\tilde{v}_t$	$\tilde{u}_t$	$\ln \frac{\tilde{V}_t}{U_t}$	$\ln \frac{\tilde{U}_t}{U_t}$	$\ln \frac{\tilde{V}_t}{J_t}$	$\ln \frac{\tilde{U}_t}{J_t}$	t	const.	R^2
(15) h_t	-0.038	-0.252*					0.441*	-0.120*					-0.0003**	8.142*	0.97
121	0.048	0.037					0.041	0.043					0.0001	0.770	
(16) f_t			0.285*						0.535*	0.151*			-0.0007*	0.099*	0.99
121			0.052						0.057	0.048			0.0002	0.035	
(17) h_t	0.179*			-0.122**			0.392*	-0.201*					-0.0006*	6.532*	0.96
101	0.045			0.051			0.056	0.059					0.0002	1.056	
(18) $\ln \frac{H_t}{J_t}$					0.324*						0.464*	0.103*	-0.0009*	0.059	0.98
101					0.046						0.065	0.037	0.0002	0.079	
(19) $\ln UE_t$	0.091	0.353*					0.274**						-0.0014*	1.426	0.89
101	0.129	0.097					0.112						0.0003	1.663	
(20) $\ln SE_t$	-0.031	-0.726*				0.454*	0.397*						-0.0002	7.169	0.98
101	0.068	0.051				0.034	0.060						0.0002	0.954	

strong negative time trend, significant at the 1% significance level, that we find in the model for UE_t also accords with a comparable finding in Gregg and Petrongolo (2005). By contrast, the model for SE_t does not seem to exhibit any trend.

In conclusion, it appears that the magnitude of the estimated negative time trend roughly halves when the flows into unemployment and vacancies are included as explanatory variables in the matching function. At the same time, the evidence in this section clearly suggests that a significant negative time trend remains. In other words, the approach based on stock-flow matching does not fully account for the time trend in standard matching functions. The next section explores a potential explanation for this remainder.

4.5 Bias from the omission of selectivity

4.5.1 Theory

It was mentioned above that the efficiency of matching might seemingly deteriorate when, for example, unemployment benefits increase and unemployed job seekers consequently raise their reservation wages, rejecting a greater share of the job offers they receive. Such changes to the unemployment benefit are exogenous to search models, but the same effect can arise endogenously from changes in labour market conditions: when unemployed workers understand that they are more sought after, they might become more picky. A tendency for workers to become more picky over time would then translate into a negative time trend. This section and the next more generally examine whether the omission of such endogenous selectivity is likely to play a role for estimated time trends.

While the following theoretical arguments apply very widely, we present them in the same set-up as Stevens (2007) (with exogenous search intensity) who highlights the role of endogenous selectivity and points to potential empirical consequences when it is omitted. Time is continuous in this set-up. For simplicity, all workers are either unemployed and searching or employed and not searching (while our empirical evaluation below will also allow for employed job seekers). Only when an unemployed worker meets a firm the productivity y of this potential match is realised as a random draw from the continuous distribution function $F(y)$ and observed by the firm and the worker. The value of search to an unemployed worker is

$$rY^U = b + \frac{m(z,V,U)}{U}\pi\left(\frac{V}{U}\right)S(y) \tag{4.28}$$

where r is the common discount rate, b is the unemployment benefit as a flow payoff, $S(y)$ is the expected surplus generated by a match, $\pi\left(V/U\right) \in (0,1)$ is the surplus share the worker obtains (e.g. through Nash bargaining with workers' bargaining power set to $\pi(V/U)$), and $m(z,V,U)$ is the matching function. Similarly, the value to a firm of offering a vacancy is

$$rY^V = -c + \frac{m(z,V,U)}{V}\left(1-\pi\left(\frac{V}{U}\right)\right)S(y) \tag{4.29}$$

where c is the flow cost of the firm's recruitment efforts. This set-up is non-standard in two ways. One is that the worker's surplus share varies with market conditions: $\pi = \pi(V/U)$. This realistic property would follow immediately if market conditions affected workers' bargaining power in a model with Nash bargaining, for example. The other non-standard feature is the dependence of the matching function on the joint reservation productivity z, so that only matches with $y \geq z$ are actually concluded. This reservation productivity is found as

$$z = rY^U + rY^V$$

In words, the match is not concluded when its productivity is so low that at least one agent prefers to continue searching. It is, however, straightforward to integrate this with standard matching functions:

$$m(z,V,U) = p(z)C(V,U) \tag{4.30}$$

where $p(z) = 1 - F(z)$ is the probability that the contact leads to a match and $C(\cdot)$

is the *contact function* which could be any matching function specification using U (or more generally job seekers) and V from models that do not distinguish between contacts and actual matches. Therefore, let us choose as before a Cobb-Douglas specification with CRS. We can then rewrite equation (4.30) in discrete time and in logarithms as follows:

$$f_t = K + \lambda \ln p_t + \beta \theta_t + \gamma t + \epsilon_t \tag{4.31}$$

If one omits, as in the case of the standard matching function, the variable $\ln p_t$, the estimate of γ' in that regression will in expectation be

$$\mathbb{E}(\hat{\gamma}') = \gamma + \lambda \tau_z \tag{4.32}$$

where τ_z is the coefficient of t in a regression of $\ln p_t$ on all other explanatory variables. While the theory in this section does not suggest a particular expression for τ_z, it does suggest the sign of the covariance between z and θ. Let us assume that firms create vacancies as long as it pays off (the free-entry condition), while the mass of workers is fixed. Then rY^V is driven to 0 so that, from equation (4.29),

$$S(y) = \frac{V}{m(z, V, U)} \frac{c}{1 - \pi\left(\frac{V}{U}\right)}$$

which allows us to substitute for $S(y)$ in equation (4.28). Noting that $z = rY^U$ when $rY^V = 0$, we thus obtain

$$z = b + \frac{V}{U} \frac{c\pi\left(\frac{V}{U}\right)}{1 - \pi\left(\frac{V}{U}\right)} \tag{4.33}$$

Hence the degree of selectivity, captured by the level of z, depends endogenously on market tightness. From equation (4.33), we can find the covariance of z and θ as

$$\begin{aligned} \operatorname{Cov}(z, \theta) &= c \cdot \operatorname{Cov}\left(e^{\theta} \frac{\pi\left(e^{\theta}\right)}{1 - \pi\left(e^{\theta}\right)}, \theta\right) \\ &= c\left[\mathbb{E}\left(\theta e^{\theta} \frac{\pi\left(e^{\theta}\right)}{1 - \pi\left(e^{\theta}\right)}\right) - \mathbb{E}\left(e^{\theta} \frac{\pi\left(e^{\theta}\right)}{1 - \pi\left(e^{\theta}\right)}\right) \mathbb{E}(\theta)\right] \end{aligned}$$

We would expect that the first derivative of $\pi(\cdot)$ is strictly positive, so that the worker obtains a higher surplus share when workers are relatively sought after. Then $\text{Cov}(z, \theta)$ will also be strictly positive. Now since $p(z)$ is falling in z, we would expect it to be falling also in θ, in the sense that the coefficient κ_z of θ_t in a regression of $\ln p_t$ on all explanatory variables is negative. With

$$\mathbb{E}(\hat{\beta}') = \beta + \lambda \kappa_z \tag{4.34}$$

this means that the estimated β will be biased downwards when selectivity is omitted.

4.5.2 Some empirical evidence

Due to data limitations, we cannot directly evaluate the effect of omitting $\ln p_t$. Instead we will exploit the relation between z and market tightness to obtain some indirect evidence. Equation (4.33) suggests that workers will be less picky when V/U is low. Hence, in a regression that omits $\ln p_t$, we would expect a higher estimated speed of matching when V/U is low. Further, we would argue that the magnitude of $\text{Cov}(z, \theta)$ is greater when higher values of θ are observed. To see this, note that

$$\mathbb{E}\left(e^{\theta} \frac{\pi\left(e^{\theta}\right)}{1-\pi\left(e^{\theta}\right)}\right) \equiv \mu$$

is a constant, so that we can write

$$\text{Cov}(z, \theta) = c\left[\mathbb{E}\left(\theta e^{\theta} \frac{\pi\left(e^{\theta}\right)}{1-\pi\left(e^{\theta}\right)}\right) - \mathbb{E}(\theta\mu)\right] \tag{4.35}$$

Suppose we add an observation of θ above the average of previous observations of θ. Since $\pi(V/U)$ will be high if θ is high, the added observation of θ is multiplied by a factor above μ in the first expectation, but is only multiplied by a factor hardly different from μ in the second expectation. The result is an increase in the magnitude of $\text{Cov}(z, \theta)$. More formally, the expressions inside the expectations in equation (4.35) are supermodular. Then the marginal effect of θ is increasing in the second factor. Because these expressions are also log-supermodular, this effect is disproportionately greater for higher θ. Therefore, when one sample contains throughout higher observations of θ than another sample, we expect the sample covariance between z and θ to be greater for the first sample. By equation (4.34) we would consequently expect the downward bias in $\hat{\beta}'$ to be larger (i.e. $|\kappa_z|$ is larger) in the first sample.

These hypotheses are testable. We thus categorise our $T = 122$ observations by the relative size of V_t/U_t: group L includes those 61 observations with market tightness in the bottom half of the ranking, and group H includes the 61 observations with market tightness in the top half of the ranking. In order to allow the estimated elasticity and speed of matching to vary between these two groups, we use a simple dummy set-up:

$$f_t = K + \beta_0 D_t^L + \beta_L \theta_t + \beta_D D_t^L \theta_t + \gamma t + \epsilon_t \tag{4.36}$$

where $D_t^L = 1$ if the market tightness in period t falls into the bottom half of the ranking, that is, if it falls into group L, and otherwise $D_t^L = 0$. The estimate of the elasticity in group H is thus recovered as $\hat{\beta}_H = \hat{\beta}_L + \hat{\beta}_D$, while $\hat{\beta}_L$ will be a direct estimate of the elasticity in group L. The results from a regression with this model are reported in line (21) of table 4.7. The estimate of β_0 is positive and significant at the 1% significance level, indicating that the speed of matching is higher in group L.[17] The estimate of β_D is also significant, suggesting a substantial difference between β_L and β_H. The fact that the estimate of β_D is positive could indeed mean that the downward bias in the estimate of β is greater when V_t/U_t is high. At the same time, the other estimated coefficients are plausible. Hence, the results in line (21) offer first evidence in support of our hypotheses.

However, these findings do not turn out to be robust. Line (22) estimates the model in equation (4.36) with all job seekers, using the same measure J_t that we obtained from our empirical model above. Here, neither the estimate of β_0 nor the estimate of β_D is significant at the 5% significance level. In line (23), the flows into unemployment and vacancies are included in the model in equation (4.36). The esti-

[17] Sedláček (2010) jointly estimates vacancies and the speed of matching. He finds a procyclical speed of matching on the U.S. labour market (albeit for an earlier time period than considered here). To confirm this finding, our empirical test would have to produce a negative and significant coefficient estimate for D_t^L because market tightness is procyclical.

Table 4.7: Matching functions with varying coefficients (OLS)

dep. var. and T	θ_t	$D_t^L \theta_t$	$\ln \frac{V_t}{J_t}$	$D_t^L \ln \frac{V_t}{J_t}$	$\ln \frac{\tilde{V}_t}{U_t}$	$\ln \frac{\tilde{U}_t}{U_t}$	$\ln \frac{\tilde{V}_t}{J_t}$	$\ln \frac{\tilde{U}_t}{J_t}$	t	D_t^L	const.	seas. dum.	adj. R^2
(21) f_t	0.621*	0.168*							-0.0015*	0.116*	-0.119*	yes	0.99
122	0.042	0.047							0.0002	0.037	0.030		
(22) $\ln \frac{H_t}{J_t}$			0.618*	0.049					-0.0016*	0.085	-0.320**	yes	0.98
101			0.042	0.040					0.0002	0.069	0.063		
(23) f_t	0.284*	0.039			0.525*	0.140*			-0.0006*	0.044	0.073	yes	0.99
121	0.055	0.042			0.061	0.049			0.0002	0.031	0.045		
(24) $\ln \frac{H_t}{J_t}$			0.309*	0.034			0.457*	0.107*	-0.0008*	0.059	0.056	yes	0.98
101			0.053	0.032			0.066	0.039	0.0002	0.055	0.076		

mates of β_0 and β_D remain insignificant, while the other coefficients appear plausible by comparison to our earlier regressions that included the inflows into unemployment and vacancies. Finally, the same conclusions also apply to a regression that includes both all job seekers and the inflows, as reported in line (24).

Overall, the results of our empirical test do not support the hypothesis that the omission of $\ln p_t$ substantively affects the estimate of β. It is in principle possible that the estimated time trend is biased by the omission of $\ln p_t$ while the estimate of β is not, but this would require that t and θ are uncorrelated. Yet the correlation between θ_t ($\ln(V_t/J_t)$) and time is -0.59 (-0.55) in our data, so that biases in the estimated time trend and in the estimate of β would probably occur together. In conclusion, the magnitude of the negative time trend that remains after the inclusion of other job seekers and flows is rather unlikely to arise from the omission of selectivity.

4.6 Conclusions

The common finding of a negative time trend in empirical matching functions has suggested that labour market performance deteriorates over time. Especially for recent years, this result is at odds with improvements that one would expect due to new information and communication technologies. Attempts to explain the time trend by real economic developments have had limited success. Therefore, we have examined the possibility that the omission of job seekers beyond the unemployed, of inflow measures, or of agents' selectivity biases the estimated coefficients in the matching function and thereby generates a negative time trend. Using recent U.S. labour market data and an empirical model that accounts for all labour market flows, we can construct data on employed and non-participating job seekers and on various flow measures. This enables us to estimate matching functions that do not omit these variables.

It turns out that the inclusion of employed and non-participating job seekers does not affect the magnitude of the estimated time trend. Yet when inflow measures are included in the matching function, this magnitude drops by about 50%. This result suggests that the finding of deteriorating labour market performance may at least partly be a statistical illusion. The remaining negative time trend is likely not explained by the omission of selectivity from the matching function: once all job seekers or the inflow measures have been included, we do not find any evidence that the estimated coefficients are substantively affected by this omission.

The actual bias from the omission of selectivity cannot be quantified here, as this would require monthly observations of workers' reservation wages. Where this rare information is available, our hypotheses may be tested more directly. Until then, it remains unclear whether the true time trend in the empirical matching function for the U.S. labour market is positive, negative, or zero, as this will depend on the size of the bias from the omission of selectivity. Indeed, it is only possible to quantify biases from the omission of other job seekers and of inflow measures because our model has allowed us to construct series for important labour market flows, on which reliable data are unavailable. Given that the constructed series seem by and large plausible and deliver estimation results in line with comparable studies, these constructed series may also prove useful for future empirical analyses.

4.A Constructed flow series

Year	Period	UE	EU	NU	UN	SE
2001	M01					
2001	M02	2131	1803	1048	654	2411
2001	M03	2389	1738	943	240	3074
2001	M04	2038	1778	1194	804	3968
2001	M05	2766	1799	902	-20	3132
2001	M06	2501	1773	1035	49	3411
2001	M07	1405	1676	966	1138	4388
2001	M08	1410	1778	1188	1097	4093
2001	M09	1630	1841	996	1107	3699
2001	M10	1784	2237	874	775	3718
2001	M11	1882	2034	1056	899	2482
2001	M12	840	1892	1145	1942	2760
2002	M01	-85	2076	968	3205	5272
2002	M02	2054	1913	1071	897	2018
2002	M03	2427	1896	1178	558	2015
2002	M04	2426	2028	904	211	3149
2002	M05	2768	1864	979	275	2666
2002	M06	2590	1730	981	127	3012
2002	M07	1652	1883	1008	1242	4092
2002	M08	1703	1911	1016	1310	3603
2002	M09	1794	1805	967	1031	3343
2002	M10	1986	1895	870	723	3270
2002	M11	2006	1859	1091	731	2325
2002	M12	787	1863	975	1931	2924
2003	M01	39	1841	1015	2937	5199
2003	M02	1936	1816	982	764	1994
2003	M03	2177	1812	1019	684	2022
2003	M04	2145	1739	1055	395	3111
2003	M05	2920	2029	959	-47	2173
2003	M06	2521	1730	1176	76	2935
2003	M07	1680	1799	921	1295	3474
2003	M08	1546	1848	953	1370	3516
2003	M09	1906	1783	921	773	3144
2003	M10	1867	1687	1023	1032	3560
2003	M11	1912	1785	871	900	2246
2003	M12	685	1525	1029	2128	3079
2004	M01	-104	1556	1129	2736	5211
2004	M02	1814	1555	870	814	2119
2004	M03	2645	1763	874	-332	2347
2004	M04	2438	1745	1035	663	3267
2004	M05	2759	1545	1178	-78	2727

2004	M06	2559	1617	1060	44	3230
2004	M07	1811	1895	925	1159	3680
2004	M08	1423	1521	1139	1383	4183
2004	M09	1866	1692	1066	955	3460
2004	M10	1988	1654	1078	610	3614
2004	M11	1936	1703	935	831	2712
2004	M12	786	1717	1051	1980	3153
2005	M01	-110	1641	994	2895	5680
2005	M02	2015	1584	1165	538	2368
2005	M03	2140	1449	1034	586	2890
2005	M04	2354	1554	1134	399	3536
2005	M05	2746	1646	1108	29	3029
2005	M06	2802	1620	1039	-16	3431
2005	M07	1811	1481	1087	875	4079
2005	M08	1427	1471	1099	1204	4587
2005	M09	1863	1791	972	692	3808
2005	M10	1670	1578	1141	1149	3825
2005	M11	2012	1601	1222	698	2688
2005	M12	618	1541	1047	2257	3229
2006	M01	227	1505	1022	2515	4894
2006	M02	1941	1465	1112	516	2580
2006	M03	2494	1676	983	277	2585
2006	M04	2306	1680	985	311	3430
2006	M05	2567	1616	929	118	3443
2006	M06	2501	1529	1176	183	3742
2006	M07	1627	1495	1226	920	4496
2006	M08	1504	1553	1050	1183	4370
2006	M09	1750	1494	1114	1102	3903
2006	M10	1548	1513	1103	1188	4148
2006	M11	1753	1518	993	613	3158
2006	M12	563	1566	1030	2143	3340
2007	M01	256	1550	1018	1974	5031
2007	M02	1898	1641	931	874	2426
2007	M03	1974	1235	1062	502	3151
2007	M04	2065	1500	954	274	3755
2007	M05	2552	1526	935	-21	3284
2007	M06	2537	1541	1008	-202	3603
2007	M07	1531	1687	813	800	4226
2007	M08	1315	1630	992	1371	4431
2007	M09	1641	1489	1086	828	3795
2007	M10	1730	1606	935	730	3922
2007	M11	1707	1582	1059	945	2880
2007	M12	559	1625	1091	1754	3057
2008	M01	310	1653	907	2261	4676
2008	M02	1776	1700	940	1076	2287

2008	M03	2268	1811	984	187	2340
2008	M04	1927	1628	862	738	3109
2008	M05	2412	1822	1448	66	2723
2008	M06	2357	1725	1034	210	3111
2008	M07	1363	1673	1183	1130	3245
2008	M08	1625	2243	1052	1134	3357
2008	M09	1573	1963	918	1240	2925
2008	M10	1656	2173	993	891	2622
2008	M11	1579	2352	957	1354	1768
2008	M12	592	2234	1032	1882	2385
2009	M01	646	2668	854	2236	3199
2009	M02	1798	2290	1109	848	1527
2009	M03	2295	2529	848	541	1343
2009	M04	2164	2373	952	705	1931
2009	M05	2610	2325	905	-158	1576
2009	M06	2251	2158	1006	649	2035
2009	M07	1827	2243	907	1436	2179
2009	M08	1540	2103	897	1170	2581
2009	M09	1960	2114	773	731	2068
2009	M10	1979	2223	1002	767	1727
2009	M11	1759	1890	877	1430	1672
2009	M12	853	1938	970	2049	1852
2010	M01	923	1803	1112	2362	2506
2010	M02	2097	2023	706	614	1036
2010	M03	2515	1790	864	56	1490
2010	M04	2413	1635	1060	87	1845
2010	M05	3182	2070	693	-165	1564
2010	M06	2346	1989	790	724	2312
2010	M07	1661	1810	1023	1128	2602
2010	M08	1472	1880	876	1072	2845
2010	M09	1990	1965	907	985	2125
2010	M10	1801	1649	1010	728	2183
2010	M11	2157	2133	691	502	1328
2010	M12	767	1690	1035	2514	2074
2011	M01	951	1681	997	2349	2638
2011	M02	1785	1558	832	795	1552

Epilogue

Admittedly, this thesis has been somewhat obsessed with paradoxes. If my Greek hasn't failed me completely, a $\pi\alpha\rho\alpha\delta o\xi o\nu$ is something contrary to one's expectation. That is, a paradox arises when theory and empirical observation don't quite fit. Which is great. After all, it's an opportunity to make actual scientific progress. Think of empirical observation as an attempt to describe reality, and think of theory as a way to imagine all sorts of causal links. If we want to know the causes of things that actually happen we'll have to make these two sides fit. Now if they don't, either a causal link wasn't sufficiently obvious for our theories to include it already, or there is a problem in the empirical methods we keep using. Hence, fixing paradoxes can really improve our theories and empirical methods.

While work on paradoxes therefore seems important, it's also rather hard. Believe it or not, economists like it when their theories and their empirical results are in harmony. When they are not, people try to do something about it. Hence, when a paradox is still there after long years, a number of attempts to resolve it have already failed or have remained somehow unsatisfactory. To then solve the paradox, you'd better have thought up an approach that doesn't fall short where previous attempts have fallen short. You'd think you just have to jump higher - more realistically, you have to run into a wall a couple of times and get a bloody nose.

So which walls have I chosen to run into? I have revisited the Diamond paradox that frictional labour markets function in practice although, in theory, the wages firms offer don't make job search worthwhile for workers. Next was the paradox that the kind of sorting we see among types in virtually every real-world matching market arises in theory only under conditions that are unlikely to hold in the real world.

And finally, we examined how it can be that declining labour market performance is a pervasive empirical finding while job search and recruiting on the internet should in theory speed up things on the labour market. None of these paradoxes is less than a decade old, and the Diamond paradox is almost a decade older than I am.

I don't claim that the solutions I could eventually propose for each of these paradoxes are somehow unique or more general than others or more elegant. But they are apparently new and they have thus far held up to scrutiny in academic presentations as well as conversations with friends and family (most of whom are not economists). This tells me that my proposed solutions are not entirely implausible. For the Diamond paradox, my solution has two parts. The first is that workers find searching worthwhile when they believe the promise in firms' adverts to pay sufficiently high wages. That much was known, but why should firms stick to their promise? Because, and this is the second part, a worker in a firm that broke its promise won't believe the firm again, and that makes the job less productive for the firm. For the next paradox, I find that the sorting we see in the real world arises under a very mild condition in theory when agents can exchange information prior to meetings, which they do in practice.

Finally, the solution I would propose for the third paradox is not in this thesis. Instead, the last chapter documents a number of clashes with the wall. But a couple of weeks after the chapter was submitted in its current form, I apparently found a solution, using still the same methods. In a new version of the chapter (available as IAB Discussion Paper 3/2012), I could present results that make the entire decline in labour market performance look like a statistical artefact, just as the chapter had hypothesised. What seems to be behind this decline is an altogether unspectacular regularity in the data on vacancies, something the current form of the chapter does not consider. The failure, in the current chapter as in virtually all papers on this topic, to account for this regularity can bias empirical results and generate a seeming decline of labour market performance where really there isn't any.

I thus do claim that the work for this thesis has ultimately produced a plausible solution for each of the three paradoxes. Time will show whether they continue to hold up to scrutiny. Even in the extreme case that none of them does, my particular crash tests might at least hold some useful information for those facing the same or similar walls later. I trust my supervisors and my examiners that the solutions I have proposed were eventually enough to merit a doctorate. While there would probably have been ways to finish this thesis much earlier (such as: keeping away from paradoxes or doing what people tell me to), finishing a thesis as such has never quite become my primary concern.

When I invest several years of my life, I hope above all that it leads to a reasonably useful contribution. After all, my memories from a now distant past suggest that I began studying economics because it analyses some of the most important and most difficult problems faced by modern societies. In Germany, unemployment had been around four million for years when I graduated from high school. Unemployment at such a scale did not only seem like a colossal waste of people's talents and experience. It also severely limited individual life aspirations, made families unhappy, and dragged down entire regions. At the same time, the numerous public debates on unemployment were all too similar to debates on football: everybody had an opinion, but those with actual expertise were far and few between. Ironically, economics was important enough to take the centre stage in numerous public debates, but also complex enough to make these debates largely fruitless. There was only one solution - it had to be studied.

So does this thesis have anything to say about unemployment? I think so. Chapter 2 considered how information helps labour markets to function at all. Its analysis implies that the way workers and firms can interact on the job might determine how reliable the firms' job adverts are. The reliability of adverts in turn affects whether unemployed workers have reason to search for jobs in the first place. Chapter 3 suggests that labour markets are improved by better information: when more precise information is exchanged before the interview, workers and firms can better prioritise the most promising interviews. Other interviews are skipped - those that would probably waste everybody's time and thereby prolong unemployment. Chapter 4 investigates why improvements from better information seemingly don't materialise. Especially in the version that has become a discussion paper, it finds that the standard method of measuring changes in labour market performance appears to be misleading. Hence it may well be that labour market performance actually improves with better information. In conclusion, as the technologies for exchanging information are developed further, I would expect that their use in labour markets and in labour market policy helps to reduce unemployment.

Berlin, April 2012

Bibliography

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

[2] Adachi, Hiroyuki (2003): "A Search Model of Two-sided Matching under Nontransferable Utility," *Journal of Economic Theory* 113, 182-198

[3] Albrecht, James W. and Axell, Bo (1984): "An Equilibrium Model of Search Unemployment," *Journal of Political Economy* 92, 824-840

[4] Anderson, Patricia M. and Burgess, Simon M. (2000): "Empirical Matching Functions: Estimation and Interpretation Using State-Level Data," *Review of Economics and Statistics* 82, 93-102

[5] Atakan, Alp E. (2006): "Assortative Matching with Explicit Search Costs," *Econometrica* 74, 667-680

[6] Autor, David H. (2001): "Wiring the Labor Market ," *Journal of Economic Perspectives* 15, 25-40

[7] Beaudry, Paul (1994): "Why an Informed Principal May Leave Rents to an Agent," *International Economic Review* 35, 821-832

[8] Becker, Gary S. (1973): "A Theory of Marriage: Part I," *Journal of Political Economy* 81, 813-846

[9] Blanchard, Olivier J. and Diamond, Peter A. (1990): "The Aggregate Matching Function," in Diamond, Peter A. (ed.): *Growth, Productivity and Unemployment*, 159-206, Cambridge, MA: MIT Press

[10] Bleakley, Hoyt and Fuhrer, Jeffrey C. (1997): "Shifts in the Beveridge Curve, Job Matching, and Labor Market Dynamics" *New England Economic Review* September/October 1997, 3-19

[11] Borowczyk-Martins, Daniel, Jolivet, Grégory, and Postel-Vinay, Fabien (2011): "Accounting for Endogenous Search Behavior in Matching Function Estimation" IZA mimeo

[12] Broersma, Lourens and van Ours, Jan (1999): "Job Searchers, Job Matches and the Elasticity of Matching," *Labour Economics* 6, 77-93

[13] Burdett, Kenneth (1978): "A Theory of Employee Job Search and Quit Rates," *American Economic Review* 68, 212-220

[14] Burdett, Kenneth and Coles, Melvyn G. (1997): "Marriage and Class," *Quarterly Journal of Economics* 112, 141-168

[15] Burdett, Kenneth, Coles, Melvyn G., and van Ours, Jan (1994): "Temporal Aggregation Bias in Stock-Flow Models," CEPR mimeo

[16] Burdett, Kenneth and Judd, Kenneth L. (1983): "Equilibrium Price Dispersion," *Econometrica* 51, 955-969

[17] Burdett, Kenneth and Mortensen, Dale T. (1998): "Wage Differentials, Employer Size, and Unemployment," *International Economic Review* 39, 257-273

[18] Burdett, Kenneth, Shi, Shouyong, and Wright, Randall (2001): "Pricing and Matching with Frictions," *Journal of Political Economy* 109, 1060-1085

[19] Burgess, Simon M. (1993): "A Model of Competition Between Unemployed and Employed Job Searchers: An Application to the Unemployment Outflow Rate in Britain," *Economic Journal* 103, 1190-1204

[20] Carrillo-Tudela, Carlos (2009): "An Equilibrium Search Model When Firms Observe Workers' Employment Status," *International Economic Review* 50, 485-506

[21] Chade, Hector (2006): "Matching with Noise and the Acceptance Curse," *Journal of Economic Theory* 129, 81-113

[22] Coles, Melvyn G. (2001): "Equilibrium Wage Dispersion, Firm Size, and Growth," *Review of Economic Dynamics* 4, 159-187

[23] Coles, Melvyn G. and Eeckhout, Jan (2003): "Indeterminacy and Di-

rected Search," *Journal of Economic Theory* 111, 265-276

[24] Coles, Melvyn G. and Petrongolo, Barbara (2008): "A Test Between Stock-Flow Matching and the Random Matching Function Approach," *International Economic Review* 49, 1113-1141

[25] Coles, Melvyn G. and Smith, Eric (1998): "Marketplaces and Matching," *International Economic Review* 39, 239-254

[26] Coles, Peter, Kushnir, Alexey, and Niederle, Muriel (2011): "Preference Signaling in Matching Markets," Stanford University mimeo

[27] Crawford, Vincent P. and Sobel, Joel (1982): "Strategic Information Transmission," *Econometrica* 50, 1431-1451

[28] Diamond, Peter A. (1971): "A Model of Price Adjustment," *Journal of Economic Theory* 3, 156-168

[29] Diamond, Peter A. (1982): "Aggregate Demand Management in Search Equilibrium," *Journal of Political Economy* 90, 881-894

[30] Eeckhout, Jan and Kircher, Philipp (2010): "Sorting and Decentralized Price Competition," *Econometrica* 78, 539-574

[31] Fahr, René and Sunde, Uwe (2004): "Occupational Job Creation: Patterns and Implications," *Oxford Economic Papers* 56, 407-435

[32] Fallick, Bruce and Fleischman, Charles A. (2004): "Employer-to-Employer Flows in the U.S. Labor Market: The Complete Picture of Gross Worker Flows," Board of Governors of the Federal Reserve System mimeo

[33] Flinn, Christopher J. (1986): "Wages and Job Mobility of Young Workers," *Journal of Political Economy* 94, S88-S110

[34] Gale, David and Shapley, Lloyd S. (1962): "College Admissions and the Stability of Marriage," *American Mathematical Monthly* 69, 9-15

[35] Gaumont, Damien, Schindler, Martin, and Wright, Randall (2006): "Alternative Theories of Wage Dispersion," *European Economic Review* 50, 831-848

[36] Goldmanis, Maris, Ray, Korok and Stuart, Erik (2009): "Until Death Do Us Part? The Marriage Model with Divorce," University of Chicago mimeo

[37] Gregg, Paul and Petrongolo, Barbara (2005): "Stock-Flow Matching and the Performance of the Labor Market," *European Economic Review* 49, 1987-2011

[38] Halko, Marja-Liisa, Kultti, Klaus, and Virrankoski, Juha (2008): "Search Direction and Wage Dispersion," *International Economic Review* 49, 111-134

[39] Hall, Robert E. and Krueger, Alan B. (2008): "Wage Formation between Newly Hired Workers and Employers: Survey Evidence," National Bureau of Economic Research mimeo

[40] Hopkins, Ed (2012): "Job Market Signalling of Relative Position, or Becker Married to Spence," *Journal of the European Economic Association* 10, 290-322

[41] Hopkins, Ed and Seymour, Robert M. (2002): "The Stability of Price Dispersion under Seller and Consumer Learning," *International Economic Review* 43, 1157-1190

[42] Hoppe, Heidrun C., Moldovanu, Benny, and Sela, Aner (2009): "The Theory of Assortative Matching Based on Costly Signals," *Review of Economic Studies* 76, 253-281

[43] Jackman, Richard, Layard, Richard and Pissarides, Christopher A. (1989): "On Vacancies," *Oxford Bulletin of Economics and Statistics* 51, 377-394

[44] Jacquet, Nicolas L. and Tan, Serene (2007): "On the Segmentation of Markets," *Journal of Political Economy* 115, 639-664

[45] Jolivet, Grégory (2009): "A Longitudinal Analysis of Search Frictions and Matching in the U.S. Labor Market," *Labour Economics* 16, 121-134

[46] Koopmans, Tjalling C. and Beckmann, Martin (1957): "Assignment Problems and the Location of Economic Activities," *Econometrica* 25, 53-76

[47] Kübler, Dorothea, Müller, Wieland, and Normann, Hans-Theo (2008): "Job Market Signaling and Screening: An Experimental Comparison," *Games and Economic Behavior* 64, 219-236

[48] Lam, David (1988): "Marriage Markets and Assortative Mating with House-

hold Public Goods: Theoretical Results and Empirical Implications," *Journal of Human Resources* 23, 462-487

[49] Lentz, Rasmus (2010): "Sorting by Search Intensity," *Journal of Economic Theory* 145, 1436-1452

[50] Lindeboom, Maarten, van Ours, Jan, and Renes, Gusta (1994): "Matching Employers and Workers: An Empirical Analysis on the Effectiveness of Search," *Oxford Economic Papers* 46, 45-67

[51] MacLeod, W. Bentley and Malcomson, James M. (1989): "Implicit Contracts, Incentive Compatibility, and Involuntary Unemployment," *Econometrica* 57, 447-480

[52] Malcomson, James M. (1999): "Individual Employment Contracts," in Ashenfelter, Orley and Card, David (eds.): *Handbook of Labor Economics*, vol. 3B, 2291-2372, Amsterdam: Elsevier Science

[53] Mare, Robert D. (1991): "Five Decades of Educational Assortative Mating," *American Sociological Review* 56, 15-32

[54] Masters, Adrian M. (2005): "Directed Search Without Wage Commitment and the Role of Labor Market Institutions," State University of New York-Albany mimeo

[55] McCall, John J. (1970): "Economics of Information and Job Search," *Quarterly Journal of Economics* 84, 113-126

[56] Menzio, Guido (2007): "A Theory of Partially Directed Search," *Journal of Political Economy* 115, 748-769

[57] Moen, Espen R. (1997): "Competitive Search Equilibrium," *Journal of Political Economy* 105, 385-411

[58] Montgomery, James D. (1991): "Equilibrium Wage Dispersion and Interindustry Wage Differentials," *Quarterly Journal of Economics* 106, 163-179

[59] Morgan, Peter B. (1998): "A Model of Search, Coordination, and Market Segmentation," University at Buffalo mimeo

[60] Mortensen, Dale T. (1970): "Job Search, the Duration of Unemployment, and the Phillips Curve," *American Economic Review* 60, 847-862

[61] Mortensen, Dale T. (1976): "Job Matching under Imperfect Information," in Ashenfelter, Orley and Blum, James (eds.): *Evaluating the Labor-Market Effects of Social Programs,* 194-232, Princeton, New Jersey: Princeton University Press

[62] Mortensen, Dale T. (2003): *Wage Dispersion: Why are Similar Workers Paid Differently?,* Cambridge: MIT Press

[63] Mortensen, Dale T. and Pissarides, Christopher A. (1999): "New Developments in Models of Search in the Labor Market," in Ashenfelter, Orley C. and Card, David (eds.): *Handbook of Labor Economics*, 2567-2627, Amsterdam: North Holland

[64] Mumford, Karen and Smith, Peter N. (1999): "The Hiring Function Reconsidered: On Closing the Circle," *Oxford Bulletin of Economics and Statistics* 61, 343-364

[65] Nagypál, Éva (2005): "On the Extent of Job-to-Job Transitions," Northwestern University mimeo

[66] Nagypál, Éva (2008): "Worker Reallocation over the Business Cycle: The Importance of Employer-to-Employer Transitions," Northwestern University mimeo

[67] Nickell, Stephen, Nunziata, Luca, and Ochel, Wolfgang (2005): "Unemployment in the OECD since the 1960s. What do we know?," *Economic Journal* 115, 1-27

[68] Peters, Michael (1991): "Ex Ante Price Offers in Matching Games Non-Steady States," *Econometrica* 59, 1425-1454

[69] Petrongolo, Barbara and Pissarides, Christopher A. (2001): "Looking into the Black Box: A Survey of the Matching Function," *Journal of Economic Literature* 39, 390-431

[70] Pissarides, Christopher A. (2000): *Equilibrium Unemployment Theory,* Second edition, Cambridge, MA: MIT Press (first edition, 1990, Oxford: Blackwell)

[71] Postel-Vinay, Fabien and Robin, Jean-Marc (2002): "The Distribution

of Earnings in an Equilibrium Search Model with State-Dependent Offers and Counteroffers," *International Economic Review* 43, 989-1016

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

[73] Ramaswami, Chitra (1983): "Equilibrium Unemployment and the Efficient Job-Finding Rate," *Journal of Labor Economics* 1, 171-196

[74] Rogerson, Richard, Shimer, Robert, and Wright, Randall (2005): "Search-Theoretic Models of the Labor Market: A Survey," *Journal of Economic Literature* 43, 959-988

[75] Roth, Alvin E. and Sotomayor, Marilda A. O. (1990): *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis,* Cambridge: Cambridge University Press

[76] Rubinstein, Ariel (1982): "Perfect Equilibrium in a Bargaining Model," *Econometrica* 50, 97-110

[77] Sedláček, Petr (2010): "Match Effciency and the Cyclical Behavior of Job Finding Rates," University of Amsterdam mimeo

[78] Shapiro, Carl and Stiglitz, Joseph E. (1984): "Equilibrium Unemployment as a Worker Discipline Device," *American Economic Review* 74, 433-444

[79] Shimer, Robert (2005a): "The Assignment of Workers to Jobs in an Economy with Coordination Frictions," *Journal of Political Economy* 113, 996-1025

[80] Shimer, Robert (2005b): "The Cyclicality of Hires, Separations, and Job-to-Job Transitions," *Federal Reserve Bank of St. Louis Review* 87, 493-507

[81] Shimer, Robert (2007): "Reassessing the Ins and Outs of Unemployment," University of Chicago mimeo

[82] Shimer, Robert and Smith, Lones (2000): "Assortative Matching and Search," *Econometrica* 68, 343-369

[83] Smith, Lones (2006): "The Marriage Model with Search Frictions," *Journal of Political Economy* 114, 1124-1146

[84] Spence, Michael (1973): "Job Market Signaling," *Quarterly Journal of Economics* 87, 355-374

[85] Stevens, Margaret (2004): “Wage-Tenure Contracts in a Frictional Labour Market: Firms’ Strategies for Recruitment and Retention,” *Review of Economic Studies* 71, 535-551

[86] Stevens, Margaret (2007): “New Microfoundations for the Aggregate Matching Function,” *International Economic Review* 48, 847-868

[87] Stigler, George J. (1961): “The Economics of Information,” *Journal of Political Economy* 69, 213-225

[88] Sunde, Uwe (2007): “Empirical Matching Functions: Searchers, Vacancies, and (Un-)biased Elasticities,” *Economica* 74, 537-560

[89] Telser, Lester G. (1978): *Economic Theory and the Core,* Chicago: University of Chicago Press

[90] van den Berg, Gerard J. and Ridder, Geert (1998): “An Empirical Equilibrium Search Model of the Labor Market,” *Econometrica* 66, 1183-1221

[91] Visschers, Ludo (2005): “Market Signaling with Frictions,” University of Pennsylvania mimeo

[92] Weiss, Andrew (1980): “Job Queues and Layoffs in Labor Markets with Flexible Wages,” *Journal of Political Economy* 88, 526-538

[93] Weitzman, Martin L. (1979): “Optimal Search for the Best Alternative,” *Econometrica* 47, 641-654