Research Statement

Ronald Wolthoff
University of Toronto

October 2018

1 Introduction

One of the fundamental goals of economic research is to understand the functioning of markets and the role they play in allocating resources. Most of my research concerns two closely related questions within this broader agenda:

1. how do market participants—e.g., workers and firms in a labor market, or buyers and sellers in a goods market or an asset market—meet and subsequently match or trade with each other?
2. how does the process that governs meetings and matches affect aggregate outcomes and the response of these outcomes to aggregate shocks?

The answers to these questions are not trivial for a number of reasons. First, market participants are generally highly heterogeneous in characteristics that affect the surplus of a potential match; e.g., workers differ in their productivity, and buyers differ in their valuation for the good that they try to obtain. Second, search and information frictions often complicate the matching process; e.g., market participants may not be able to (easily) observe each other's locations, characteristics or actions. Third, overcoming those frictions may involve various decisions and costs; e.g., firms have to choose how much time and how many resources to dedicate to the screening of applicants. Finally, both the search for a trading partner and the match that is formed may last for some time, creating dynamic considerations; e.g. workers incur the cost of job search before receiving the returns associated with finding a job, the magnitude of which in turn depends on the length of the employment spell.

My research in this area focuses on two types of markets in particular, labor markets and goods markets. My work on both markets is primarily theoretical, with the nuance that the papers on labor markets are somewhat more applied and sometimes include empirical work, while the papers on goods markets tend to be a bit more fundamental. I discuss the individual papers in section 2 and 3, respectively.[1] Besides this work, I have a few papers on other questions regarding the labor market, which I discuss in section 4.

2 Search and Matching in Labor Markets

2.1 Background

The literature on *search and matching* in the labor market was pioneered by the 2010 laureates for the Nobel Memorial Prize in Economic Sciences, Peter Diamond, Dale Mortensen, and Christopher Pissarides. Models

[1] I omit references throughout and refer to the literature reviews in my papers for the appropriate credit to prior work. Various surveys of the literature are available, e.g., Petrongolo and Pissarides (2001), Rogerson et al. (2005), Nobel Committee (2010), Albrecht (2011), Peters (2013), Chade et al. (2017) and Wright et al. (2017).

following their framework have traditionally abstracted from the details of the interaction between workers and firms leading up to a potential match. Instead, these models have generally assumed that bilateral meetings between unemployed workers and firms with a vacancy occur randomly at a rate determined by a matching function which depends only on the total number of vacancies and the total number of unemployed.

While this approach has been very fruitful and has generated many new insights, the true matching process in the labor market may of course not be invariant to policy changes and may not be stable over time or over the business cycle. In recent years, various authors have therefore started to develop micro-foundations for this process by explicitly modeling the strategic interaction between workers and firms that leads to matching. Some of this work maintains the assumption that search is random and that wages are determined and/or revealed ex post. However, other work in this strand of literature assumes that search is directed. That is, firms post wages—or more generally, contracts—which workers observe ex ante. This allows workers to target applications to particular firms, taking into account that they will face more competition from other workers (in expectation) at higher wages or more attractive contracts.

2.2 Contributions

In the steady state of continuous-time search models, whether random or directed, transitions between employment states are assumed to take place uniformly across time. In **"It's About Time: Implications of the Period Length in an Equilibrium Search Model"** (*International Economic Review*, 2014), I use French administrative data on the begin and end dates of unemployment spells to document that this assumption does not hold in reality. Instead, transitions exhibit clustering: employment relationships are much more likely to start or end on the first day of the month than on any other day. A similar pattern can be observed within a week, with transitions on Monday being most frequent.

To understand how this clustering affects labor market outcomes, I develop a search model in which workers randomly receive wage offers in continuous time, but transitions between employment states are restricted to certain pre-determined time points. The frequency of these time points is determined by a parameter, the *period length*. Infinitesimally short periods result in convergence to a continuous-time model in which meetings between workers and firms are bilateral and transitions take place uniformly across time. However, when the time period is longer, workers may receive multiple job offers per period and transitions are clustered.

I calibrate the model to US data to analyze the effect of the period length on aggregate outcomes. I find that, if the underlying meeting rates are constant, then an increase in the period length leads to a higher unemployment rate (because more job offers are now turned down). In addition, this increase in period length increases the average wage in the economy (because firms have to compete more aggressively for any worker that they meet). The variance of cross-sectional wages is however non-monotonic in the period length. Hence, the period length and the corresponding bilateral or multilateral nature of meetings are not innocuous choices and warrant careful consideration.

Although allowing for richer meeting patterns is a step in the right direction, the meeting process is of course not exogenously determined in reality, but (at least partially) the outcome of workers' and firms' endogenous decisions. For example, workers can send either more or fewer applications per unit of time. Pieter Gautier, Jose Luis Moraga-Gonzalez and I explore this idea in **"Search Costs and Efficiency: Do Unemployed Workers Search Enough?"** (*European Economic Review*, 2016) by developing a discrete-time version of a random-search model in which workers differ in their search cost and—based on this cost—choose how

many applications they want to send per period. Upon meeting an applicant, firms make a wage offer without knowing how many other wage offers this applicant receives. This information friction causes firms to randomize their wage offers, yielding frictional wage dispersion, which in turn provides workers with an extra incentive to send multiple applications.

We use Dutch data on workers' search intensities to identify (bounds on) the distribution of search costs, which we subsequently use to solve the problem of a social planner in this environment. We find that some workers search too little while other workers search too much relative to the efficient levels. On the one hand, too many workers do not apply at all, failing to internalize the firms' part of the match surplus. On the other hand, other workers send too many applications in an attempt to get a high wage offer, not taking into account that this rent-seeking behavior makes it harder for other workers to match. Interestingly, an increase in unemployment benefits or the minimum wage could improve efficiency along both margins, since both changes increase the expected return from a first application but decrease the marginal gain from subsequent applications.

Whereas the previous paper abstracts from productivity heterogeneity, it is the main focus in **"Simultaneous Search with Heterogeneous Firms and Ex Post Competition"** (*Labour Economics*, 2009), in which Pieter Gautier and I analyze how workers allocate their applications across two sectors that differ in productivity. In the model, search is random within each sector. Each firm initially offers an applicant his outside option, but can increase its wage offer if the applicant has multiple offers. We then prove that workers sending two applications will never spread these applications across the two sectors in equilibrium. Intuitively, a worker's wage in this environment is determined by his second-best offer, which is the same when he sends both applications to the low-productivity sector as when he diversifies his applications across sectors. Since the probability of receiving a job offer is higher in the low-productivity sector, the former option is more attractive. Workers will therefore randomize between applying twice to the high sector and applying twice to the low sector, if the productivity differences are not too large. However, this application behavior is inefficient. Rather than one worker sending two applications to the high sector and another worker applying twice in the low sector, a social planner prefers both to diversify their applications so as to create scope for two matches in the high sector rather than one. We prove that this inefficiency survives for many parameters values even if search is directed rather than random within each sector. This reveals that the source of the inefficiency is novel in the literature.

In all the above papers, as in much of the literature, the firm side of the model is very stylized in the sense that—although each worker can meet multiple firms—firms meet at most one worker at a time. In other words, firms face no real recruitment decisions, such as how many applicants to interview. This assumption is typically made to keep the model tractable, as a complicated network of meetings arises when both sides of the market are in contact with multiple potential trading partners. Nevertheless, studying firms' recruiting decisions is important because they create an additional intensive margin to the matching process, which—as recent empirical work has suggested—can affect the response of the labor market to aggregate shocks.

I analyze this point in detail in **"Applications and Interviews: Firms' Recruiting Decisions in a Frictional Labor Market"** (*The Review of Economic Studies*, 2018). In particular, I develop a directed-search model in which workers choose how many applications to send and firms choose how many applicants to interview. In the model, a firm's applicants differ along two dimensions: i) applicants differ in their match-specific productivity, which is revealed during an interview; ii) applicants differ in their outside options, i.e., some applicants may turn down the firm's offer because of a better offer somewhere else. Both aspects imply

that firms generally interview multiple applicants.

In this environment, I show that aggregate outcomes—such as uniqueness of equilibrium and the cyclicality of firms' recruiting decisions—crucially depend on firms' recruiting cost and workers' search cost. Calibration of the model to the US labor market indicates that these costs are relatively small but nevertheless important in the sense that they give rise to substantially different outcomes than an economy without intensive margins. First, the calibrated values for these costs imply existence of a continuum of equilibria, differing in payoffs, output and efficiency, rather than a unique equilibrium. Second, given selection of a particular equilibrium, these costs generate a matching process that responds to productivity shocks in a way that is more in line with the data.

A natural next question is how firms' recruitment decisions interact with workers' search behavior when workers differ in their productivity ex ante. I address this question with Benjamin Lester in **"Interviews and the Assignment of Workers to Jobs"** (in progress, 2016). Ex ante productivity differences complicate the problem considerably, so we analyze the model under the standard assumption that workers send one application per period. A second difference compared to the previous paper is that firms interview their applicants sequentially, until the cost of an extra interview exceeds the marginal gain. We prove that the planner's solution in this environment can be decentralized by firms posting a menu of wages—one for each type of worker—as well as information about the productivity level at which they will stop interviewing more candidates. First simulation results suggest that the impact of aggregate shocks is largest exactly in the case that has received least attention in the literature, i.e., when firms interview multiple but not all applicants.

Even though the above papers and other work in this area indicate that the meeting process matters for labor market outcomes, empirical evidence on how workers and firms exactly meet and match with each other is still very scarce. In **"Opening the Black Box of the Matching Function: The Power of Words"** (revise and resubmit at *Journal of Labor Economics*, 2016), Ioana Marinescu and I try to shed some light on the first stages of this process using data from job ads on the employment website CareerBuilder.com. We document that only 20% of the job ads contains information about the wage. At first sight, this may seem to undermine the empirical relevance of directed-search models of the labor market. However, we argue that this conclusion is not justified for two reasons.

First, we show that much information about the job is conveyed through the job title, i.e., the firms' own description of the position that appears at the top of the ad. These job titles are much more detailed than standard occupational classifications, distinguishing between different levels of seniority (e.g., junior versus senior accountant) or areas of specialization (e.g, c++ versus java programmers). As a result, the wage dispersion within a job title is much smaller than the wage dispersion researchers have documented for occupations, which means that workers may have a reasonably good idea about their potential earnings from the job title alone and do not require explicit wage announcements.

Second, we find that—everything else equal—the relation between the wage and the number of applicants is positive (among firms that post wages), which is in line with the predictions of directed-search models but contradicts earlier findings in the literature. We show that the fact that we account for job titles explains the difference; if one omits job titles and only controls for occupations, one finds a spurious negative relationship, because higher-paying job titles within an occupation tend to get fewer applicants.

3 Competing Mechanisms in Goods Markets

3.1 Background

Decentralized trade between a large number of agents does not only take place in the labor market, but also in the goods market. The process by which agents meet and subsequently trade in these markets is studied in the literature on *competing mechanisms*. A typical environment in this literature is one in which buyers have independent private valuations for the good and sellers choose mechanisms to induce revelation of buyers' types. The interaction between agents closely resembles directed-search models of the labor market: sellers publicly post their mechanisms, which are observed by buyers before they decide which seller to visit, and both types of agents face a trade-off between the probability of trade and the payoff conditional on this event. The standard assumption is that while each buyer meets only one seller, each seller can meet multiple buyers. In particular, meetings are often modeled with an urn-ball technology. This technology specifies that a buyer chooses randomly among all sellers that post the buyer's preferred mechanism. The result is that the number of buyers at a seller follows a Poisson distribution with mean dependent on the seller's mechanism.[2] In equilibrium, sellers then post auctions, but without additional instruments—such as reserve prices—that are often used in monopolistic settings. As a result, the equilibrium is constrained efficient.

3.2 Contributions

In **"Meeting Technologies and Optimal Trading Mechanisms in Competitive Search"** (*Journal of Economic Theory*, 2015), Benjamin Lester, Ludo Visschers and I study a version of the standard environment in which buyers learn their valuations upon meeting a seller. We analyze why decentralization of the planner's solution in this environment requires only an auction and no payments by or to the buyers that do not trade. After all, constrained efficiency places strict requirements on two distinct margins. First, from an ex ante point of view, it requires the proper allocation of buyers to sellers. Second, from an ex post point of view, it requires the proper allocation of the good once buyers arrive at a seller. It is well-known that an auction is the right instrument to satisfy the latter requirement, but a priori there is no reason to expect that this mechanism is sufficient to satisfy the former requirement as well.

We therefore extend the environment to allow for a wide class of alternative meeting technologies that include both the possibility that a seller may meet at most 4 buyers or may only be able to meet a buyer when at least 7 buyers show up. We prove that—for any meeting technology—the equilibrium mechanism consists of an auction and an appropriately chosen meeting fee, to be paid by (or to, when negative) all buyers meeting the seller. This meeting fee exactly prices the externalities created by the meeting technology. That is, if a buyer visiting a seller makes it harder (easier) for that seller to meet other buyers, then the meeting fee is positive (negative). We establish that there is no externality—i.e., the fee equals zero—if and only if the meeting technology satisfies a novel condition, *invariance*. We demonstrate that the urn-ball technology satisfies invariance, which explains why the existing literature did not identify meeting fees as a necessary element of the equilibrium mechanism. However, the class of invariant technologies is not restricted to urn-ball, as we show by characterizing another micro-founded example that satisfies the relevant condition.

In **"Meetings and Mechanisms"** (in progress, 2017), Xiaoming Cai, Pieter Gautier and I consider a similar environment, except that buyers know their types before choosing which seller to visit. This small change in assumption complicates the problem considerably, because agents no longer necessarily all attempt to trade

[2] While some papers in this literature deal with finite numbers of agents, which yields a binomial distribution of buyers at each seller, I restrict attention to large markets with a continuum of buyers and sellers.

in the same (sub)market; we prove that the planner's solution requires that the number of submarkets be less than or equal to one plus the number of buyer types. We make a methodological contribution by introducing a one-to-one transformation of meeting technologies—i.e., the probability that a seller meets at least one buyer of a certain type—which considerably simplifies the analysis. Sellers again post auctions combined with a meeting fee, but the queue of buyers that they attract is not necessarily uniquely determined since buyers are ex ante heterogeneous. We derive a number of weak conditions on the meeting technology that guarantee uniqueness of the queue. In this case, there exists a unique equilibrium, which we show to be constrained efficient. Finally, we extend the model to allow for heterogeneity in seller's own valuations and derive conditions that support assortative matching, i.e., high-value sellers attract high-value buyers.

Xiaoming Cai, Pieter Gautier and I build on these insights in a closely related paper, titled **"Search Frictions, Competing Mechanisms and Optimal Market Segmentation"** (*Journal of Economic Theory*, 2017). In this paper, we consider the inverse problem, which is about identification: what we can learn about the meeting process from equilibrium outcomes such as the degree of market segmentation? We answer this question by deriving necessary and sufficient conditions on the meeting technology for the two outcomes that have received most attention in the literature: i) perfect separation, i.e., a separate submarket for each type of buyer, and ii) perfect pooling, i.e., a single market with all buyers and sellers. We establish that perfect separation arises for any distribution of buyers' valuations if and only if meetings between buyers and sellers are always bilateral. In that case, auctions do not generate any revenue and the meeting fees essentially act like posted prices. In contrast, the auction is crucial under perfect pooling, which arises if and only if the meeting technology satisfies a novel condition, *joint concavity*. We conclude the paper with a detailed discussion of this condition—e.g., we show that invariant technologies, such as urn-ball, satisfy joint concavity, while technologies with negative externalities between meetings do not.

Following the literature, the above three papers all take the meeting technology as exogenously given. However, as in labor markets, it seems likely that meetings in real-life goods markets are in fact endogenously determined: meeting potential trading partners is often costly—e.g., due to transportation costs or the opportunity cost of time—and this cost determines the number of meetings chosen by buyers and/or sellers. Jointly with Benjamin Lester and Ludo Visschers, I explore the implications of this idea in a paper titled **"Competing with Asking Prices"** (*Theoretical Economics*, 2017).

In this paper, buyers again differ in their valuation for a seller's good, but they need to meet the seller in order to learn this valuation and this meeting is costly to them. We find that this substantially changes the mechanism that sellers post in equilibrium. In particular, we show that auctions are no longer optimal, precisely because they are so good at extracting surplus. That is, a buyer who observes a seller posting an auction will conclude that this seller is not a very attractive trading partner because he does not sufficiently compensate him for incurring the inspection cost. As result, this seller will attract relatively few buyers, leaving him with a low payoff after all.

Instead, we establish that sellers maximize their revenue—and decentralize the planner's solution—by posting what we call an *asking price* mechanism. Such a mechanism features sequential meetings between the seller and potential buyers, which stop when a buyer submits a bid equal to the posted asking price, triggering immediate trade. If no such bids are received and the seller has no more potential buyers, he recalls the highest bidder, as long as that bid exceeded his own valuation. The mechanism does not include any other instruments, such as meeting fees or reservation prices, making it similar to the way goods are sold in many real-life markets.

4 Other Work on Labor Markets

Three of my papers concern the labor market but do not deal with the meeting process between workers and firms. In **"Becker meets Ricardo: Multisector Matching with Communication and Cognitive Skills"** (*Journal of Law, Economics and Organization*, 2015), Robert McCann, Xianwen Shi, Aloysius Siow and I study how heterogeneous individuals match in a competitive education and labor market. In the model, individuals live for two periods and are born with heterogeneous communication skills and cognitive ability. When young, individuals acquire cognitive skills in schools, based on their own cognitive ability and the cognitive skill of their teacher. As adults, individuals choose whether to become a worker, a manager, or a teacher. Workers and managers match within firms, where managers with higher communication skill manage more workers. Similarly, teachers with higher communication skill teach more students.

Although the model features matching between individuals with multidimensional skills across multiple sectors, we show that it remains very tractable. The equilibrium in the model decentralizes the planner's solution, which can be obtained by solving a linear-programming problem. In equilibrium, adults with high cognitive skills and high communication skills become managers or teachers, while the remaining adults become workers. Interestingly, individuals who expect to become managers or teachers later in life will invest discretely more in education than marginally different individuals who expect to become workers. We use simulation to demonstrate that the model is capable of generating a positive correlation between cognitive and communication skills as well as a long right tail in the earnings distribution, even if initial abilities are independent and uniform.

Some of the more technical issues that arise in this environment as a result of the infinite-dimensional nature of the planner's linear program problem are addressed in **"Academic Wages and Pyramid Schemes: A Mathematical Model"** (*Journal of Functional Analysis*, 2015) with Alice Erlinger, Robert McCann, Xianwen Shi, and Aloysius Siow. In particular, this paper analyzes whether the wages of the individuals with the highest cognitive skill—who become the teachers of future teachers of future teachers of ..., and so forth—can become unbounded. We find that this is not the case, although the gradient of wages at the top becomes unbounded if the product of a teacher's number of students and his relative contribution to their cognitive skill is larger than one.

Finally, in **"Early Retirement Behaviour in the Netherlands: Evidence from a Policy Reform"** (*De Economist – Netherlands Economic Review*, 2010), Rob Euwals, Daniel van Vuuren and I study how financial incentives affect workers' retirement behavior. Historically, early retirement programs in the Netherlands were very generous but not actuarially fair in the sense that they provided strong financial incentives to retire at exactly the age of first eligibility, as benefits could not be carried forward. In the early 1990s, Dutch labour unions and employer organizations agreed to transform these programs, making them less generous but more actuarially fair by making the pension benefit increase in the retirement age. The transition towards the new programs created variation in individuals' incentives to postpone retirement by date of birth and sector of industry. We exploit this variation to estimate hazard rate models for the retirement age on a large administrative dataset from the Netherlands. We find that both the income effect, due to lower retirement wealth, and the substitution effect, due to lower implicit taxes on postponing retirement, have induced workers to retire later, with the latter effect being more substantial.

References

Albrecht, J. W. (2011). Search theory: The 2010 Nobel Memorial Prize in Economic Sciences. *Scandinavian Journal of Economics*, 113(2):237–259.

Chade, H., Eeckhout, J., and Smith, L. (2017). Sorting through search and matching models in economics. *Journal of Economic Literature*, 55(2):493–544.

Nobel Committee (2010). Markets with search frictions: Scientific background on the Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel 2010.

Peters, M. (2013). *The Handbook of Market Design*, chapter Competing Mechanisms. Oxford University Press.

Petrongolo, B. and Pissarides, C. A. (2001). Looking into the black box: A survey of the matching function. *Journal of Economic Literature*, 39:390–431.

Rogerson, R., Shimer, R., and Wright, R. (2005). Search-theoretic models of the labor market: A survey. *Journal of Economic Literature*, 43(4):959–988.

Wright, R., Kircher, P., Julien, B., and Guerrieri, V. (2017). Directed search: A guided tour. mimeo.