Bilateral Uncertainty in a Model of Job-Market Screening with Intermediaries*

Paul Schweinzer
Birkbeck College, University of London
P.Schweinzer@econ.bbk.ac.uk

Abstract

We look at a job-market model of bilateral uncertainty. Workers are uncertain about what job descriptions advertised by firms really mean and firms are uncertain about the qualifications of workers before they are interviewed. Both types of uncertainty can be resolved but both processes are costly. Intermediaries (recruiters) can perform the job matching but only at the cost of transforming the firm's objective. Our model is one in which potential employees are screened competitively through the offer of pairs of ⟨*incentives*, *job descriptions*⟩. What drives the results of the model is that (*i*) job descriptions are fundamentally incomplete and therefore costly to communicate and (*ii*) that the firm has inferior knowledge of the (endogenous) workers' inside options than recruiters have. Therefore it may be cheaper for the firm to hire an intermediary than to negotiate by itself. We study the role of an external intermediary who uses a Vickrey auction to discriminate between workers' qualifications. We find that the presence of a recruiter changes the structure of the offered packages by distorting the job descriptions upwards. We show that the presence of competition among recruiters alters the market outcome and lowers the social efficiency of matching if workers choose recruiters randomly. If we allow recruiters to announce their specialisation, however, workers self-select and full efficiency can be achieved. (JEL *D44, D82, E24, J41.* Keywords: *Bilateral uncertainty, incomplete contracts, screening, job-market auctions.*)

Introduction

If you take a look at the job offers sections in papers like the *Financial Times* or the *Economist*, you will realise that the jobs one finds advertised there ask for an extraordinarily high level of qualification even for reasonably low-profile jobs. Of course, it is not very likely that a person with such high qualifications would actually accept these jobs.[1] Why put up a smokescreen by asking for exorbitant qualifications? Our model explains this through the presence of an incompleteness in the job descriptions that firms draw up in order to attract potential employees.

* Thanks to *Pix Sandhu* and *Arup Daripa* (3rd January 2002).

[1] A *London School of Economics* job advertisement (from 1/12/00) for the position of an assistant to the Director says "the successful applicant will have an outstanding degree in the social sciences and some experience of working in a large and complex organisation. They will be academically oriented but administratively aware, energetic, innovative, diplomatic and a networker [...] the post will be fixed term to 30 April 2004; the salary will be within the range of 21,616 - 27,347 inc." If you find this irresistible, contact LSE-Resource-Centre@lse.ac.uk.

The same incompleteness gives recruiters a role to play in the job-market. We show that their objective is furthered by keeping job descriptions vague and providing them with an upward drift towards more complex jobs.

Our model incorporates bilateral uncertainty in the job-market.[2] Firms are uncertain about the aptitude of the workers they plan to hire, and workers are uncertain about their future responsibilities. In order to resolve this uncertainty, firms screen their prospective employees by offering a certain menu of job descriptions and wages that only those (privately informed) workers who are willing to perform the specified type of work for the wage offered will accept. These jobs, however, are not completely specified in the job advertisements and therefore the screening procedure will not necessarily be effective without a clarification of what the advertisement actually means. The fuzzy job descriptions in advertisements may, however, establish a lower and upper bound on workers' expectations about the required qualification.

The workers, on the other hand, face a job advertisement that only incompletely describes the job they are supposed to do. In order to ascertain the offered contract, workers need to participate in an interview—that is, they have to engage in a costly activity in order to gather information about jobs. The interview is costly for both the worker and the firm but this expenditure is necessary to resolve the worker's uncertainty about the job and to enable the firm to screen.

It is potentially costly for a highly qualified worker to sign a contract before the contract incompleteness is resolved because as soon as the contract is signed, all property rights are with the firm.[3] This means that ex-post renegotiation will drive down the worker's wage to whatever her qualification is worth to the firm. Before the interview takes place the worker does not know the exact contract details—therefore higher qualified workers will show up if the fuzzy job specifications are distorted upwards but not if they consider them tailored for a lower qualification-type than their own.

The bilateral and possibly repeated expenditures for resolving the contract incompleteness on the one hand and the firm's uncertainty about the workers' qualifications on the other hand are the reasons for frictions in the market. Job-market intermediaries such as recruiters, or head hunters, may take advantage of their comparative advantage which lies in the fact that they work for a number of different principals and therefore get information on a wider range of job/worker types than would be possible for a single firm at any given time. In particular, intermediaries gain superior knowledge on workers' inside options, i.e. their alternative job offers. This additional information puts the recruiter in a better position to negotiate contracts

[2] The literature is not consistent in this terminology. We use the term bilateral or two-sided uncertainty for a setting where all involved parties are mutually uncertain about elements of the other's objective—we do not imply that there are only two parties to the contract.

[3] Tirole (1999) defines a property right to be a bundle of decision rights that determine the use of an asset. As soon as the worker signs an employment contract, the firm has complete control over both its assets and the labour-time of the worker. Therefore everything that is not determined through the employment contract ('residual' property rights) can be ex-post filled in to the advantage of the firm. The classic definition of the role of control rights in incomplete contracts is due to Grossman and Hart (1986) and Hart and Moore (1988) who apply the idea exclusively to physical assets (and not employment contracts).

than firms. The recruiter uses these endogenous inside options as informational background for a second-price sealed-bid auction that induces the revelation of the participating workers' qualification-types.

Our model is neither applicable to jobs that have virtually no skills requirements (the contract being complete) nor to placements that are at least partially determined by political considerations such as CEO placements (the contracts being *very* incomplete).

The literature in relation to our model falls into four broad categories: the literature on bilateral uncertainty, job-market hiring, the use of auctions in this process, and the question of whether recruiters optimally specialise on the market place or not. The literature on bilateral (or two-sided) uncertainty was initiated in a bargaining context by Chatterjee and Samuelson (1983) and extended and surveyed by Cramton (1992) who uses delay as a signal of the players' valuations. The most recent related result is due to Wang (2000) who shows that regardless of the bilateral uncertainty, trade will always occur when there is a surplus to be shared and delay-signalling is possible. These models do not discuss the role of an intermediary and restrict themselves to bargaining situations. Under these conditions, the basic determinant of the outcome is the players' individual discount rate while we analyse a more specific job-market situation where the information asymmetry is much more severe. There seems to be no closely related work that includes intermediaries in the literature on job-market hiring. Recent general surveys were presented by Autor (2001) and Fryer (2001). The idea to use a double auction for job-market screening appears in Hall and Lazear (1984) but is not discussed further in that paper. Another contribution is Brown (1980) who develops a model of a labour market income insurance that is tested empirically. The main issue is an investigation into the differences of the realised market wages from the predicted marginal value products and the main finding is that the labour market outcome differs from an auction only in the short run. The idea of a labour market auction is at the centre of the analysis of Julien, Kennes, and King (2000) who use an auction to derive a matching function in a macroeconomic context. Their focus is not on the analysis of the microeconomics of the screening procedure itself but on the aggregated outcome of negotiations with many players on both sides. From their dynamic analysis they derive a matching function that endogenously generates the properties that are usually taken to be exogenous in labour market matching models. This particular literature has been recently surveyed by Pissarides and Petrongolo (2001). The topic of spatial competition has been given considerable attention in the industrial organisations literature. Its principal findings have been surveyed by Graitson (1982) and Bester, de Palma, Leininger, Thomas, and von Thadden (1996). Our full separation of recruiters result is due to a modification on Hotelling's original model that—while keeping the marketplace linear—allows an analysis in the spirit of Salop (1979). Although our recruiter's objective is a sum and not a product as in d'Aspremont, Gabszewicz, and Thisse (1979), their analysis of maximum vs. minimum differentiation along a linear city is in principle applicable to our model and agrees with our result.

The paper is organised as follows. Section 1 sets out the basic model and its assumptions.

Our main result is characterised in section 2 in a screening setting where we assume that a single job negotiation involves only one recruiter. This assumption is justified in section 3 where we show that recruiters specialise in a linear city à la Hotelling. We conclude with a discussion of the social desirability of such a specialisation.

1 The model

In the classical screening model of Rothschild and Stiglitz (1976) the task-level t is usually interpreted as the speed of the production line or the number of units produced per time interval. We use a wider interpretation of t as a (potentially multi-dimensional) description of a particular job that forms—together with a wage specification—the basis of an employment contract. Both elements $\langle w, t\rangle$ of a contract offer convey information about the job the firm is advertising. But in an incomplete job description, either element may be fuzzy, misleading or missing entirely. Therefore to the uniformed worker, the offered contract will be a region rather than a precise point. To distinguish the region $\mathbf{t}$ from the point t we use boldface notation.[4]

1.1 Workers

There are $I > 1$ risk-neutral workers with ability types $\theta_i \in [\theta_l, \theta_h]$ $(0 < \theta_l < \theta_h < \infty)$, that are privately known to the workers but not to the hiring or recruiting firms. However, after a worker is actually hired, her type becomes (ex-post) visible to the firm because the firm can monitor her effort perfectly well and relate it to her realised output. The worker's type is interpreted as the productivity level that she is capable of achieving. The proportion of workers with productivity θ_i or lower is given by the distribution function $F[\theta_i]$ which is assumed to be strictly increasing, atomless and publicly known. Moreover we assume workers' types to be drawn from this distribution in a statistically independent fashion.

The idea that it is ex-ante costly for a worker to gather information on a job is due to Stigler (1961). In our model this means that attending an interview and thereby obtaining complete information on the contract costs the worker $C(t)$. The information a worker holds about the job offered could be thought of as being gradually refined through successive interviews but as a first approach we assume that all uncertainty is completely resolved after only one interview—the region $\langle[\underline{w_o}, \overline{w_o}], \mathbf{t}\rangle$ shrinks to the point $\langle w, t\rangle$.

The worker's outside option u_θ^0 is the utility she can obtain by not participating in the labour market. It is nonnegative and linear in her productivity type minus a fixed cost of being unemployed that is normalised to zero

$$u_\theta^0 \equiv \alpha\theta, \quad \alpha \in [0, 1]. \tag{1.1}$$

[4] If we take the advertisements' wage offer to be 'hard', the offered region is a line segment. Usually, however, firms give an interval for their wage offers which gives an interior to our region.

This outside option is known perfectly well to the worker but, being based on her type, not at all to the firm or recruiter.

Definition 1. *The worker's* ***inside option*** *is defined as the utility she can obtain from an alternative contract* $\langle w_x, t_x \rangle$.

Inside options are private to the participants in the job negotiations.

The worker suffers an unobservable cost $c(t, \theta)$ from fulfilling the specified task t that depends on the worker's ability θ. In order to ensure a well-behaved solution we assume $c(\cdot)$ to be twice continuously differentiable and the Spence-Mirrlees ('single-crossing') condition to hold for all $t > 0$

$$c(0, \theta) = 0, \quad c_t(t, \theta) > 0, \quad c_{tt}(t, \theta) > 0, \quad c_\theta(t, \theta) < 0, \quad c_{t\theta}(t, \theta) < 0. \tag{1.2}$$

We assume that the task level t has a direct influence β on the worker's productivity

$$\theta(1 + \beta t), \quad \beta \in [0, \infty). \tag{1.3}$$

From (1.2) we know that higher types are more productive than lower types and from the previous condition we know that this difference is weakly increasing as we go up the qualifications scale. The worker's utility (depending on wages w, her job t and her own type θ) is linear

$$u_\theta(w, t, \theta) = w(t) - c(t, \theta) - C(t). \tag{1.4}$$

We assume that if a worker is indifferent about two offered contracts, she always accepts the one with the higher wage. If she is indifferent about working or not she accepts the offered contract.

1.2 Firms

There are $J > 1$ identical firms that competitively hire workers. All share the same constant returns to scale technology and are risk-neutral. We normalise the price of the firms' output to 1. Firms offer employment contracts that consist of incentives and job descriptions $\langle w_o, t_o \rangle$ and hire everyone who accepts this contract. But the specified jobs are fuzzy in both dimensions: The firm is only able to offer a range of wages $[\underline{w_o}, \overline{w_o}]$ and a range of tasks. The scale of job descriptions is an ordinal ranking identifying jobs asking for increasing qualifications. There is no 'natural' correspondence between any single job description t_i and a particular ability-type θ_i but it is sensible to assume that the highest vacant t should be filled by the highest available θ.

The firm's linear utility from a type θ worker in a job t is

$$u_f(w, t, \theta) = \theta(1 + \beta t) - w(t) \tag{1.5}$$

which it will maximise subject to participation of the worker and/or the recruiter. Hence the wage a type-θ worker will be offered for doing a job t under perfect competition is

$$w(\theta, t) = \theta(1 + \beta t). \tag{1.6}$$

In addition, the firm faces a cost $C(t)$ of communicating the job descriptions as incompletely specified by the advertisement to other parties.

Firms accept contracts proposed to them by the recruiter if the job is in the offered region (in a sense we will make more precise shortly) and the utility they get from the proposed contract is at least as high as the utility they would get were they to negotiate the contract themselves.

1.3 Contracts

Contracts are pairs of incentives and job descriptions $\langle w, t \rangle$. We restrict attention to contracts of the form $\langle w \in W, t \in T \rangle$ where we assume $t \in T \subseteq \mathbb{R}$ to be simply an ordinal ranking that identifies particular jobs and $w \in W \subseteq \mathbb{R}$ to be just a monetary wage.

We assume that ex-ante the workers hold beliefs about the expected wage offer $w_o \in [\overline{w_o}, \underline{w_o}]$. Apart from identifying a particular job, the job description is also used as part of a screening device. A firm can at the most sign one contract per worker and a worker can accept at most one contract. The firm can, however, offer a menu of contracts to screen the workers.

Contracts are incomplete in the sense that the job descriptions they contain do not fully specify the job that the worker is supposed to do.[5] The reason for this incompleteness is that firms are unable to completely specify ex-ante which contingencies will determine its operations and its internal division of labour in the future. Even if this were possible, the cost of enforcing such a contract would be prohibitive.[6] A consequence of this incomplete job specification is that it is costly to communicate to a worker what is expected of her (the same applies if the contract is negotiated by a recruiter). We take the gap

$$2\nu(t) \equiv \|sup(\mathbf{t}) - inf(\mathbf{t})\| \subseteq T \tag{1.7}$$

between the extremes of task levels that the worker could reasonably expect the advertisement t to ask for to be a measure of the extent of incompleteness of a particular job description ($\nu(t) = 0$ is a complete contract). We assume this incompleteness $\nu(t)$ to be weakly increasing

[5] Tirole (1999) defines an incomplete contract as one that merely makes "assumptions on a class of feasible mechanisms" instead of a precise delineation of "feasible outcomes by incentive constraints" that would be done were a complete contract possible.

[6] Maskin and Tirole (1999a) point out that this is not an entirely satisfactorily explanation: Consequently, they call an incomplete contract such as our job descriptions *apparently* incomplete. Their critique implies that we could alternatively assume t to be private to the company. At this point, however, we consider the approach of using a contractual incompleteness that is reduced through interviews the more natural and use the 'unforeseen contingencies' framework as developed by Maskin and Tirole (1999b) and Al-Najjar, Anderlini, and Felli (2001) to justify our use of incomplete contracts.

in t. That is, we assume that it is more difficult to precisely describe a complex job than a simple job.[7]

The firm's uncertainty about the worker's outside options is partially resolved by the worker showing up at the interview advertised by the offer $\langle w_o, t_o \rangle$. By participating, the worker sends the signal which is interpreted as her 'gross'-reservation wage being in the interval $w_o \pm \varepsilon(w_o, t_o)$.[8] Hence we interpret $\varepsilon(w,t)$ to be a measure of the firm's uncertainty about the workers outside option which we assume to be weakly increasing in t.

1.4 Recruiter

There are $K \geq 1$ recruiters who can sign contracts with some workers and some firms and refer the former to the latter. If a worker who was referred on by a recruiter and a firm sign a contract $\langle w, t \rangle$, then the recruiter gets the one-off payment $M_w > 0$.

The recruiter first gets a job description sufficiently communicated by a firm to be able to effectively refer workers back and then communicates this job description to some workers. After these external interviews, the types of those workers who the recruiter actually refers to the firm are suitable for the jobs the firm wants to fill—that is, we assume recruiters to be effective. This is implemented by a fixed fee $\gamma > M(w)$ that the recruiter has to pay if he refers a type θ_i to the firm whose outside option $u^0_{\theta_i}$ is higher than the utility $u_{\theta_i}(w_i, t_i)$ that this worker can possibly attain in the offered range of tasks.

This corresponds to the θ_i worker's individual rationality constraint

$$(ir)_w : \qquad u_{\theta_i}(w_i, t_i) \geq u^0_{\theta_i}, \qquad \forall t_i \in t_o \pm \nu(t_o). \tag{1.8}$$

Like all parties, the recruiter faces an increasing cost $C(t)$ of communicating job descriptions to workers. The interpretation is that it is increasingly expensive to communicate complex job descriptions.

The recruiter's position in the negotiation process is better than the firm's because he has access to and information on a number of other firms and workers in the course of his normal business activity. Therefore he has superior information on the market situation (e.g. the current number of vacancies, the employment opportunities for certain type-ranges) that gives him information on the workers' inside options that would vanish if the recruiter were integrated into the personnel department of the firm. Since the recruiter knows what type of jobs are open to a particular job candidate at any time, he knows the worker's inside options perfectly well. The firm lacks this knowledge.

The recruiter's objective is given by

$$u_r(w,t) = M(w) - 2C(t). \tag{1.9}$$

[7] Notice that since we construct the job descriptions on an ordinal scale, our measure of incompleteness ν is just a ranking as well. This is innocuous because we only use it to compare alternatives to *one* offered job.

[8] We discuss the signalling-mechanism in more detail in section 2.

1.5 Players and timing

There are $I + J + K$ players in the recruiting game. Firms and recruiters are risk-neutral, workers are risk-averse. We are looking at a game with the following timing:

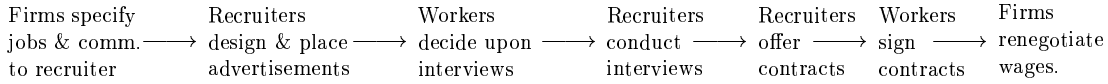

Figure 1: The timing of the recruiting game.

2 Individual contract negotiation

We can distinguish four distinct informational settings: (type certainty / complete contracts) where we can achieve first-best, (type uncertainty / complete contracts) and (type certainty / incomplete contracts) which resemble two-sided principal-agent problems because there is no scope for a recruiter and (type uncertainty / incomplete contracts) which we focus on in the following. In this setting of bilateral uncertainty the firm decides for every $\nu(t)$ of contract incompleteness whether it is optimal to resolve the contract incompleteness and its uncertainty about the worker's inside options by hiring a recruiter at cost $M(w)$ or to negotiate by itself.

Theorem 1. *In any Bayesian Nash equilibrium of the recruiting game with unobservable worker types and incomplete contracts advertising $\langle w_o, t_o \rangle$, the contract that can be negotiated by the recruiter generates a lower cost to the firm than the contract which it can negotiate on its own.*

Proof. We denote the type of an arbitrary worker who arrives at the interview by $\hat{\theta}$.

1. **Firm**. The firm knows the interval of the contract incompleteness $2\nu(t_o) > 0$ and the range of wages it specified to accommodate unknown market conditions.[9] Therefore the firm constructs the region of possible interpretations of the offered contract as the dashed box of the left panel of figure 2. The firm can determine the extreme types θ'_l, θ'_h whose outside options are compatible with this offered contract region:

$$\begin{aligned} u_{\theta'_l}(\underline{w_o}, t_o - \nu) &= u^0_{\theta'_l}, \\ u_{\theta'_h}(\overline{w_o}, t_o + \nu) &= u^0_{\theta'_h}. \end{aligned} \tag{2.1}$$

   Thus the firm knows the two limiting cases of the worst possible type accepting the worst specified contract and the best possible type accepting the best specified contract. The distance between their outside options is defined as the interval of the firm's uncertainty

[9] Nothing changes if we reverse this interpretation: The firm specifies a range of wages *because* it is uncertain about the workers' outside options.

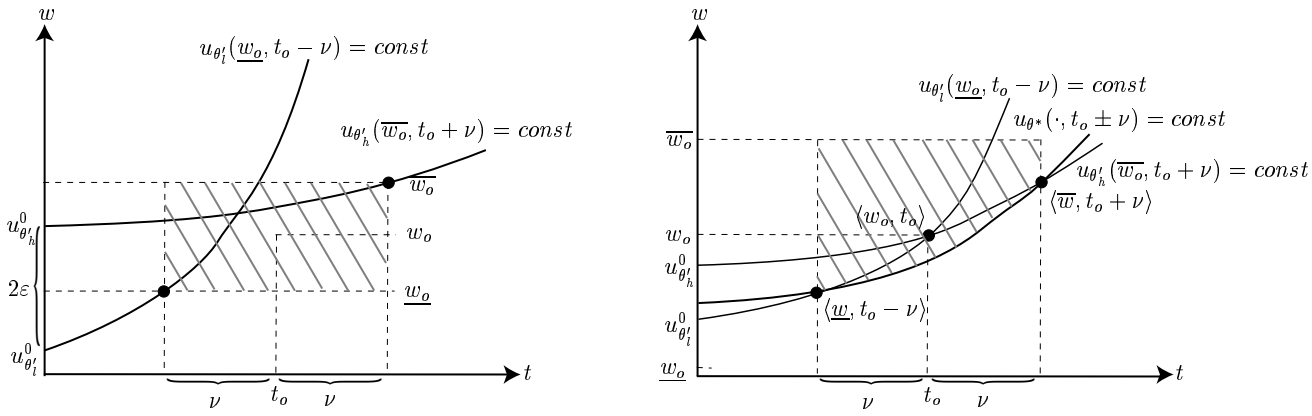

Figure 2: (Left) the firm's uncertainty about the worker's outside options and its ignorance of inside options. (Right) the recruiter's uncertainty about the worker's inside options.

about the workers' outside options

$$2\varepsilon(w_o, t_o) \equiv [u^0_{\theta'_l}, u^0_{\theta'_h}]. \tag{2.2}$$

2. **Recruiter**. The recruiter can identify the same two extreme types θ'_l, θ'_h that he knows have outside options that make them at least indifferent to accepting the offered job $\langle w_o, t_o \rangle$. But he can go further: The lower type θ'_l is indifferent between the offered job and an hypothetical inside option at $t_o - \nu$. We call this hypothetical contract $(\underline{w}, t_o - \nu)$. Similarly, the highest possible type compatible with the offered job θ'_h identifies a hypothetical inside option $(\overline{w}, t_o + \nu)$. From the Spence-Mirrlees condition (1.2) there is only one (again hypothetical) worker's type whose indifference curve passes through both $(\underline{w}, t_o - \nu)$ and $(\overline{w}, t_o + \nu)$.[10] This is the recruiter's optimal type θ^* he designs his contract offer for.[11] This type is the highest possible type who does not reject some contracts in the specified region with positive probability because she finds her outside options more tempting. At the same time this is the lowest type who is capable of doing both the jobs $\langle \underline{w_o}, t_o - \nu \rangle$ and $\langle \overline{w_o}, t_o + \nu \rangle$ without rejecting some 'reasonable' jobs with positive probability.

Now the recruiter makes use of his knowledge of the inside options he can offer to the

[10] To see this let $1 = \langle w_1, t_1 \rangle, 2 = \langle w_2, t_2 \rangle$, and $3 = \langle w_3, t_3 \rangle$ with $t_1 < t_2 < t_3$. Assume further there are three workers with $\theta_l < \theta_m < \theta_h$. Let θ_l be indifferent between 1 and 2, θ_m be indifferent between 1 and 3, and θ_h be indifferent between 2 and 3. From $c_t(\cdot) > 0$ we know that $w_1 < w_2 < w_3$. Consider now a type $\tilde{\theta} < \theta_m$ who is indifferent between $u^0_{\tilde{\theta}}$ and 1. Since $c_{\tilde{\theta}}(\cdot) > c_{\theta_m}(\cdot)$ for all $t > 0$, the curve cannot intersect once more with θ_m's indifference curve. Hence there is no other curve fulfilling the Spence-Mirrlees conditions going through both 1 and 3.

[11] Since $t_o - \nu < t_o < t_o + \nu$ we know that the wages are $\underline{w} < w_o < \overline{w}$. It is easy to see that the type θ^* worker will always prefer the contract $\langle w_o, t_o \rangle$ to both $(\underline{w}, t_o - \nu)$ and $(\overline{w}, t_o + \nu)$. In fact we have shown this result in the previous footnote with $\langle \underline{w}, t_o - \nu \rangle = 1, \langle w_o, t_o \rangle = 2, \langle \overline{w}, t_o + \nu \rangle = 3$.

θ^* worker. Since he can rank these options along the job description axis he can identify the next lower inside option. We call this contract $\langle w_{oo}, t_{oo}\rangle$. The recruiter computes the optimal worker's outside option as determined by market conditions as

$$u^i_{\theta^*} \equiv max[u^0_{\theta^*}, u_{\theta^*}(w_{oo}, t_{oo})]. \tag{2.3}$$

The recruiter has now determined both the slope and intersect of the optimal worker's level set. Being threatened by the punishment $\gamma > M(w)$, he will accept the recruitment contract only if

$$(u_{\theta^*}(\underline{w}, t_o - \nu) = u_{\theta^*}(\overline{w}, t_o + \nu)) \geq u^0_{\theta^*}. \tag{2.4}$$

The recruiter's exploitation of his knowledge of worker's inside options is illustrated in the left panel of figure 3. We define $\overline{w}$ to be the wage the recruiter offers to the optimal type θ^* accepting the initially offered contract, this is the highest possible offer the recruiter can justify ex-post to the firm. Therefore $\overline{w}$ is the highest possible offer he will make.

We conclude that

$$\overline{w_o} > \overline{w} \tag{2.5}$$

because otherwise the firm would have specified the possible range of wages $[\underline{w_o}, \overline{w_o}]$ in a way that will not allow her to hire the worker of optimal type. She therefore would have misspecified the employment contract and the firm has no hope of finding a suitable worker under the prevailing market conditions. If and only if (2.5) is true, the recruiter will accept the recruitment contract.

3. **Worker**. The candidate worker $\hat{\theta}$ knows that in all contracts negotiated by the recruiter (2.5) holds. Moreover, the worker knows that the firm knows the prevailing market conditions no better than to allow for a wage offer $\overline{w_o}$. She will therefore turn down any contract offered by a firm that specifies a lower wage than $\overline{w_o}$. She will, however, accept an offer by a recruiter involving $\overline{w} \geq w_o$ if she prefers contract

$$u_{\hat{\theta}}(w_o, t_o) \geq u_{\hat{\theta}}(w_{oo}, t_{oo}) \geq u^i_{\theta^*}. \tag{2.6}$$

Therefore the theorem holds for every properly specified contract fulfilling (2.5). Since the firm knows that in equilibrium she will not find an employee if the recruiter turns down her offered employment contract, she will not offer such contracts but modify the advertisements.

□

Therefore the firm will use a recruiter for its hiring purposes. We proceed to discuss which contract in the dashed region in the right panel of figure 2 is likely to be negotiated by a recruiter.

Example: We look at three workers $i \in \{1, 2, 3\}$ with ability $\theta_l' < \theta_1 < \theta^* < \theta_2 < \theta_3 < \theta_h'$ that all would accept contract $\langle w_o, t_o \rangle$. Let's assume a setting where θ_1 will accept $\langle \overline{w}, t_o - \nu \rangle$ but turn down $\langle \underline{w}, t_o + \nu \rangle$. Similarly, let θ_2 accept $\langle \underline{w}, t_o + \nu \rangle$ but turn down $\langle \overline{w}, t_o - \nu \rangle$, respectively. θ_3 is even better that θ_2. The recruiter knows that the firm's best offer is $\langle \overline{w_o}, t_o + \nu \rangle$ and he will stay under this wage offer because it was reached using the highest possible θ_h. Let's assume that the next highest inside option is outside the productivity reach of all three workers. This situation is illustrated in the right-hand panel of figure 3.

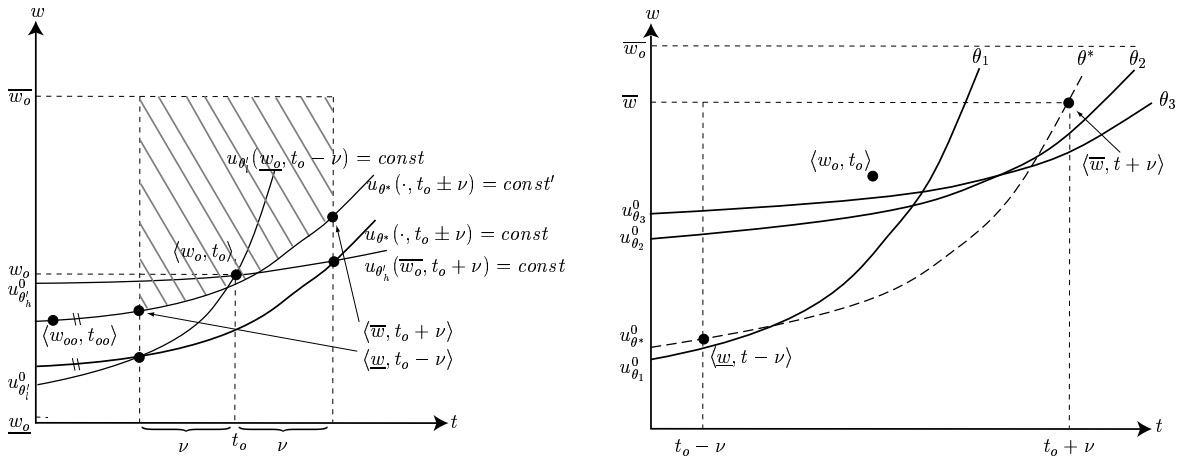

Figure 3: (Left) The inside option $\langle w_{oo}, t_{oo} \rangle$. (Right) the Vickrey auction setup.

In this example there are only two types who prefer the full range of the incomplete contract to $\langle w > \overline{w}, t_o + \nu \rangle$.[12] Regardless of which worker he chooses, the recruiter maximises his utility at complexity $t_o + \nu$. Therefore the recruiter will offer a contract $\langle \cdot, t_o + \nu \rangle$ ∘

Theorem 2. *In any Bayesian Nash equilibrium of the recruiting game with unobservable worker types and an incomplete contract advertising t_o, the implemented contract $\langle w^*, t^* \rangle$ exhibits*

$$(t^* = t_o + \nu) > t_o. \tag{2.7}$$

Proof. We know from $c_t(\cdot) > 0$ that along any worker's indifference curve, the contracted wage will be maximised at the interval-border $t_o + \nu$. Any other task level will be suboptimal for *both* the recruiter and the worker. Therefore the two share the same objective but have different constraints. The worker's constraint is to ensure participation (ir_w) and the recruiter's

[12] For a worker to prefer the full range of the incomplete contract to $\langle w, \cdot \rangle$ means that at wage w, she does all jobs described in the contract.

constraint is to be ex-post accountable (ac_r):

$$\begin{aligned} \arg\max_{\langle w,t\rangle} u_r(w,t) = M(w(t)) - 2C(t) &\geq u_r^0 \\ \text{s.t.}(ir)_w: \; w(t) - c(t,\hat{\theta}) - C(t) &\geq u_{\hat{\theta}}^0 \\ (ac)_{r1}: \; t &\in [t_o - \nu, t_o + \nu] \\ (ac)_{r2}: \; w &\in [\underline{w_o}, \overline{w_o}]. \end{aligned} \tag{2.8}$$

The first order condition of this problem gives $t^* = t_o + \nu$ as the optimal task level. The intuition is that since $M(w)$ is increasing, the recruiter will pick the point where w is largest. □

The recruiter's remaining problem is to distinguish between all workers willing to accept contracts $\langle w^* \leq \overline{w}, t^* = t_o + \nu \rangle$.

Example: We proceed in the setting of the previous example and consider a two person second-price sealed-bid auction with private values θ over the employment contract between the two remaining competitors θ_2 and θ_3 at complexity $t_o + \nu$.[13] We interpret the wage-demands w_i, $i \in \{2,3\}$ workers put forward as bids. The recruiter announces his value to be $\overline{w}$—if bids stay above this wage, then no contract is signed. The interviews are conducted separately and without an opportunity for the candidates to communicate. Candidates submit a single bid. The candidate with the lowest bid wins and 'pays' the second lowest bid.

It is well known that under the set of assumptions detailed in theorem 2 such an auction results in the contract going to the better qualified type θ_3 at the wage demanded by the lower qualified worker θ_2. The auction allows the recruiter to calculate the optimal wage offer w^* if he has some knowledge of the distribution of workers' types $F[\theta]$. Since we assume this distribution to be publicly known, w^* is the wage part of the contract the recruiter offers ₀

The firm cannot herself use an auction to screen workers because she is ignorant about the worker's inside options. Therefore she would not know how to interpret bids from workers. Since workers know this, the firm's uncertainty will be fully exploited and their lowest acceptable wage-offer remains at $\overline{w_o}$.

We know from the revenue equivalence theorem that this result holds for a wide class of auctions—therefore the type of auction we choose is of no great consequence. We generalise the above example to the following theorem.

Theorem 3. *In any Bayesian Nash equilibrium of the recruiting game with the recruiter announcing his value $\overline{w}$, unobservable worker types and an incomplete contract advertising t_o, the worker who accepts is the highest qualified worker available for this contract. She will sign at the indifference wage w^* of the second best qualified competitor for the job.*[14]

Proof. The setting described is an $I+1$-player, second-price sealed-bid (Vickrey) auction with private values. We summarise our assumptions:

[13] If the workers were not ex-post perfectly well informed about the offered contract through the interview, the appropriate setting would be a double auction as suggested in Hall and Lazear (1984).

[14] If there is only one job candidate, the implemented wage is the recruiter's value $\overline{w}$.

1. The recruiter only auctions off one contract at a time.
2. No worker wants more than one job.
3. The workers types $\hat{\theta}$ are drawn independently from the interval $[\theta_l, \theta_h]$. From the revelation principle we know that we can contend ourselves with a direct mechanism based only on the type instead of bidding functions $w(\theta, t_o + \nu)$.
4. The number of job candidates $\hat{\theta}_i$ is finite.
5. The distribution $F[\hat{\theta}]$ is atomless and $F'(\hat{\theta}) > 0$.
6. All participants are risk-neutral.
7. No worker expects a utility higher than her inside option from not being offered the job.
8. The recruiter's value of $\overline{w}$ is known to all participants.

Myerson (1981) shows that the winner of this auction is the candidate with the highest type. The wage she is 'paying' is the second highest qualified worker's wage demand. □

Remark: In principle, the firm could get more information about workers' types by offering a menu of contracts. She could then screen a worker against her own inside options that she is well-informed about. But this does *not* solve its problem of not being able to distinguish between the workers' types. Therefore this strategy is not effective in our setting.

3 Social desirability of intermediaries

The preceding discussion shows that it is *constrained efficient* for the firm to use the services of an intermediary. In order to discuss the social implications of the above theorems, we add some terminology. Let Θ represent the pool of unemployed workers and T the set of open job descriptions.

Definition 2. *A job match (θ_i, t_i) is called **socially efficient** if the highest qualified worker with $\theta_i \in \Theta$ is assigned to the most complex vacant job $t_i \in T$. That is, a job description t_i is filled socially efficiently by θ_i iff*

$$|\{t \in T | t \geq t_i\}| = |\{\theta \in \Theta | \theta \geq \theta_i\}|. \tag{3.1}$$

We call a job match *strictly* efficient if the above inequalities are strict. We now look at the individual matching behaviour of one recruiter h among competing recruiters. We call the set of jobs available to this recruiter T^h and the pool of workers he can refer Θ^h.

Definition 3. *A job match (θ_i^h, t_i^h) is called **individually efficient** for recruiter h if the highest qualified worker who is a client of h with $\theta_i^h \in \Theta^h$ is assigned to the most complex vacant job $t_i^h \in T^h$ that the recruiter is hired to fill. That is, a job description t_i^h is filled individually efficiently with θ_i^h through recruiter h iff*

$$|\{t^h \in T^h | t^h \geq t_i^h\}| = |\{\theta^h \in \Theta^h | \theta^h \geq \theta_i^h\}|. \tag{3.2}$$

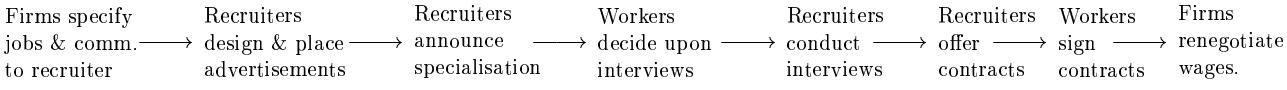

Figure 4: Modified timing—the third signalling step is inserted—of the recruiting game.

Lemma 1. *Every recruiter fills his jobs individually efficiently.*

Proof. The recruiter starts with the contract calling for the highest qualification because that contract gives him the highest expected profit. Since the recruiter fills each job with the highest qualified worker available, he will allocate each worker efficiently to a suitable job description. □

An implication is that if there is only one recruiter in the job-market and he refers all jobs, then the matching is socially efficient.

Theorem 4. *If there are several recruiters in the job-market competing for job referrals, then the matching is not socially efficient.*

Proof. Assume there are two recruiters A and B who specialise in the upper and lower half of the economy's available jobs, respectively. Assume a continuum of worker-types are drawn independently from the uniform distribution: A will place his clients individually efficiently in the upper half of the qualifications scale while B will place exclusively and individually efficiently in the lower qualifications segment. In terms of expectations, no job match in the economy is strictly socially efficient. □

3.1 Recruiter specialisation

What types of workers are likely to be attracted by recruiters if they specialise their placement activities? To discuss this question, we consider a linear job-marketplace à la Hotelling of 'length' $[0, 1]$ on which jobs described by t are uniformly distributed. For a start let us look at 2 recruiters A, B who have the same characteristics as in the above proof. They announce their average specialisation by choosing the corresponding $\hat{t} \in [0, 1]$. Thus we change our previous timing to the structure of figure 4.

We look at a continuous, uniformly distributed population of workers identified by their type $\theta \in [\theta_l, \theta_h]$ that we take to map without loss of generality into $[0, 1]$. We assume that workers who apply at a recruiter who has a widely differing specialisation face a linear disutility given by the extent of their type 'mismatch' $w \equiv \|\tau(\theta) - \upsilon(t)\|$. The interpretation is that this distance is proportional to the likelihood of expending time at interviews but not finding a suitable job. Recruiters announce their specialisation according to the maximisation of expected profits. This is illustrated in figure 5 below. It is easy to see that the unique solution to this simple uniform two-recruiter problem with $\theta_l = 0, \theta_h = 1$ is for *both* recruiters to locate at $t^* = 1$. This clearly represents a pooling equilibrium where workers cannot self-select and the inefficiencies discussed

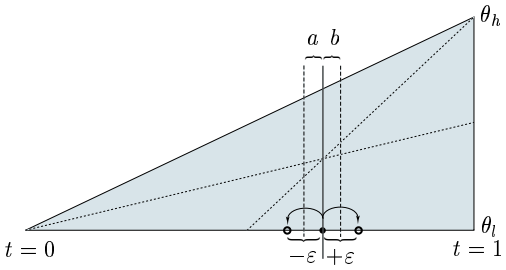

Figure 5: 2 recruiters A, B do not specialise on the linear job-market.

above apply. Fortunately, the extension to the case of many recruiters is not straightforward; especially the strategic differences between peripheral recruiters that have only one competitor and recruiters that face competition from two rivals (at the interior of the linear job-market) is problematic. To resolve this, we assume that both peripheral positions $t = \{0, 1\}$ are held by exogenous (e.g. state) institutions.[15] We sum up the assumptions on our model of product differentiation à la Hotelling as follows:

Assumption 1. *The linear job-market consists of an increasing scale of job descriptions* $t \in [0, 1] \equiv T$. *There are* n *recruiters announcing their specialisation* t_i, $i \in \{1, \ldots, n\}$ *who place a continuum of workers with types drawn from* $\theta \in [0, 1]$ *at jobs described by* t.[16] *Workers face a danger of not being placed by a recruiter that is proportional to* $|\theta - t|$.

Assumption 2. *Firms maximise expected profits proportional to the worker's types. For simplicity we take the expected profit per worker to be* ιt *for positive* ι.

Assumption 3. *As pointed out in section 1, firms and recruiters face a cost* $C(t)$ *of communicating job descriptions from firms and to workers. For simplicity we assume* $C(t) = \kappa t$ *for positive* κ.[17]

Assumption 4. *An exogenous authority subsidises the two peripheral recruiters by half the average profits of recruiters.*

Theorem 5. *Given the above assumptions, a continuum of workers represented in their job search by* n *specialised recruiters will achieve full separation on the linear job-marketplace.*

Proof. A single recruiter $i = 1$ announcing his specialisation $\hat{t}_1$ maximises expected profits proportional to

$$\Pi(\hat{t}_1, \cdot) = \int_{\frac{\hat{t}_1}{2}}^{\frac{1-\hat{t}_1}{2}} \iota t \, \mathrm{d}t - \int_{\frac{\hat{t}_1}{2}}^{\frac{1-\hat{t}_1}{2}} \kappa |t - \hat{t}_1| \, \mathrm{d}t \tag{3.3}$$

[15] It is relatively easy to justify this for $t = 0$. There is insufficient representation of workers with (nearly) no qualifications because expected profits for the recruiter are very low. It is more difficult to argue for a state agency that takes away the highest qualified workers. Or is it? Shouldn't universities keep their best students for future training and research? Wouldn't it be nice to have the people with the highest ability enter the highest public offices?

[16] Nothing precludes the recruiters' specialisations from being identical.

[17] An alternative interpretation in our model is that recruiters first need to achieve an informational advantage over firms in order to be able to sell their product. This necessitates some ex-ante expenditure.

which simplifies to

$$\iota\left(\frac{1}{8}+\frac{1}{4}\right)(1-\hat{t}_1)+\kappa\left(\frac{3\hat{t}_1}{4}-\frac{5\hat{t}_1^2}{4}-\frac{1}{8}\right) \tag{3.4}$$

which has a unique interior maximum at

$$\hat{t}_1^*=\frac{3\kappa-\iota}{10\kappa}. \tag{3.5}$$

For suitable parameters κ and ι, this strategy leads the single recruiter to a profit maximum at the interior of the linear job-market he claims. For the simplest case of $\kappa=\iota=1$, this maximum is at $\hat{t}_1^*=\frac{1}{5}$. Exactly the same argument applies (for modified interval borders and values) for a general number of recruiters n. Hence recruiters fully separate.[18] □

Notice that in our analysis in the previous sections we made use of this full separation result. The good news implied by the above result is that although we cannot hope to achieve efficiency without coordination in general, we obtain full separation and thereby full efficiency as the outcome of a modified Hotelling game that reflects workers' ability to self-select on a signal by the recruiters.

Conclusion

We provide an environment for the study of bilateral uncertainty in a screening model of the job-market. Our model thus represents an extension to the existing literature on bilateral moral hazard. The central conclusion the paper establishes is that if intermediaries find it worthwhile to participate in the job-market, they will distort upwards the descriptions of jobs in advertisements and they will systematically refer more highly qualified personnel to the firms than those would hire by themselves. The reason for this is a fundamental information advantage on inside options the recruiter enjoys by working as a common agent for several principals. The auction model we employ gives a clear prediction of the optimal wage offer a job offer should contain. We indicate that, in principle, competition among recruiters lowers the social efficiency of job matching. We end on a happier note by showing that as soon as we allow for workers to self-select upon a signal of the recruiters' specialisation, we can achieve full separation and thereby restore efficiency.

References

Al-Najjar, N., L. Anderlini, and L. Felli (2001): "Unforeseen contingencies," Unpublished mimeo.

Autor, D. H. (2001): "Wiring the labour market," *Journal of Economic Perspectives*, 15, 25–40.

Bester, H., A. de Palma, W. Leininger, J. Thomas, and E. von Thadden (1996): "A Noncooperative Analysis of Hotelling's Location Game," *Games and Economic Behavior*, 12(2), 165–86.

[18] For a general version of this proof see Schweinzer (2001).

Brown, J. N. (1980): "How close to an Auction is the Labor Market?," Princeton working paper 134.

Chatterjee, K., and W. Samuelson (1983): "Bargaining under Incomplete Information," *Operations Research*, 31, 835–51.

Cramton, P. (1992): "Strategic delay in bargaining with two-sided uncertainty," *Review of Economic Studies*, 59, 205–25.

d'Aspremont, C., J. J. Gabszewicz, and J.-F. Thisse (1979): "On Hotelling's 'Stability in Competition'," *Econometrica*, 47, 1145–50.

Fryer, R. G. (2001): "Economists' models of discrimination: An analytical survey," *Unpublished manuscript, University of Chicago.*

Graitson, D. (1982): "Spatial competition *a la* Hotelling: A selective survey," *The Journal of Industrial Organisation*, 31(1/2), 13–25.

Grossman, S., and O. Hart (1986): "The Costs and Benefits of Ownership: A Theory of Lateral and Vertical Integration," *Journal of Political Economy*, 94, 341–68.

Hall, R., and E. Lazear (1984): "The excess sensitivity of layoffs and quits to demands," *Journal of Labor Economics*, 2, 233–57.

Hart, O., and J. Moore (1988): "Incomplete contracts and renegotiation," *Econometrica*, 56, 755–85.

Julien, B., J. Kennes, and I. King (2000): "Bidding for Labor," *Review of Economic Dynamics*, 3, 619–49.

Maskin, E., and J. Tirole (1999a): "Two remarks on the property rights literature," *Review of Economic Studies*, 66, 139–50.

——— (1999b): "Unforeseen contingencies and incomplete contracts," *Review of Economic Studies*, 66, 83–114.

Myerson, R. B. (1981): "Optimal Auction Design," *Mathematics of Operations Research*, 6(1), 58–73.

Pissarides, C., and B. Petrongolo (2001): "Looking into the Black Box: A Survey of the Matching Function," *Journal of Economic Literature*, 38, 390–431.

Rothschild, M., and J. Stiglitz (1976): "Equilibrium in competitive insurance markets: An essay on the economics of imperfect information," *Quarterly Journal of Economics*, 80, 629–49.

Salop, S. C. (1979): "Monopolistic competition with outside goods," *Bell Journal of Economics*, 10, 141–56.

Schweinzer, P. (2001): "A note on the exixtence of a unique symmetric pure strategy equilibrium in a modified version of Hotelling's location game," *Unpublished manuscript, Birkbeck College.*

Stigler, G. (1961): "The economics of information," *Journal of Political Economy*, 69, 213–25.

TIROLE, J. (1999): "Incomplete Contracts: Where do we stand?," *Econometrica*, 67(4), 741–81.

WANG, R. (2000): "Separating equilibria in a continuous-time bargaining model with two-sided uncertainty," *International Journal of Game Theory*, 29(2), 229–40.