Contracting with persuasive agents

Benjamin Davies*

August 2, 2022

Abstract

I study a contracting game played by a recruitment agent and their client. The client pays the agent if they find a match. Search costs depend on market thickness, which the agent observes but the client does not. The agent can persuade the client to pay more by manipulating their beliefs about market thickness and prospective match qualities. Persuasion benefits the agent more when they can shirk than when competition prevents them from shirking.

1 Introduction

Workers and firms may struggle to find each other due to search frictions. However, both parties can overcome these frictions by hiring recruitment agents, who specialize in finding worker-firm matches. This paper analyzes agents' incentives to manipulate their clients' beliefs about market thickness and prospective match qualities.

In Section 2, I describe a contracting game in which a client (e.g., a worker or firm) pays an agent to find a match. Search costs depend on market thickness, indexed by a binary state that the agent observes but the client does not. The client has correct probabilistic beliefs about market thickness. These beliefs determine how much the client pays the agent for finding a match under the optimal contract. The client offers a low payment if they hold sufficiently strong beliefs that the market is thick and true search costs are low. The threshold belief at which the client is indifferent between low and high payments depends on the agent's cost structure, the mean match quality in thin markets, and the agent and client's patience.

Knowing the optimal contract depends on their client's beliefs, the agent manipulates those beliefs to obtain a contract they prefer. In Section 3, I derive the optimal signals for the agent to send the client about market thickness and match qualities. The client is a rational Bayesian and

*Department of Economics, Stanford University; bldavies@stanford.edu. This paper was written as a class project for ECON 285 ("Matching and Market Design") at Stanford in Fall 2021.

knows the agent wants to persuade them, but can be persuaded nonetheless. Persuasion benefits the agent more when the market is less likely to be thick and when the mean match quality in a thin market is higher.

When describing the contracting game, I assume competitive forces prevent the agent from shirking when the client offers a payment sufficient to cover effort costs. I relax this assumption in Section 4, replacing it with the client being able to find matches on their own. Persuasion benefits the agent more when they can shirk than when competition prevents them from shirking. I discuss other extensions of my analysis in Section 5.

Throughout this paper I use language specific to labor markets. However, my analysis generalizes to other matching markets in which people employ (human or machine) agents to overcome search frictions. For example, real estate agents, who match home buyers and sellers, can exploit their informational advantage and manipulate clients' beliefs for financial gain (see, e.g., Arnold, 1992; Levitt and Syverson, 2008). Likewise, dating app providers have incentives to make users believe there are many desirable matches to be found on their app. More generally, my analysis applies in any context where a relative expert (e.g., a matchmaker with superior information about market conditions) benefits from providing services and, thus, has incentives to persuade clients those services are more valuable than they are.

1.1 Related literature

This paper relates to the literature on delegated search (see, e.g., Postl, 2004; Lewis and Ottaviani, 2008; Lewis, 2012; Ulbricht, 2016; Zorc et al., 2019).[1] That literature focuses on determining when delegated search is preferable to self-search, and optimal incentive schemes when delegates' effort or results are unobserved. In contrast, I focus on how the delegate can manipulate the delegator's beliefs to make them more likely to delegate and offer higher rewards.

This paper also relates to the literature on information asymmetries in matching markets. That literature studies notions of stability when information is incomplete (e.g., Bikhchandani, 2017; Chen and Hu, 2020; Liu et al., 2014) and mechanisms for resolving incompleteness (e.g., Artemov, 2021; Ashlagi et al., 2020; Kanoria and Saban, 2021). I study how market participants interact with the matching mechanism itself—here, a recruitment agent—and that mechanism's incentives to reveal information about potential matches.

Finally, this paper lies at the (seemingly small) intersection of the literatures on information design and contract theory. Those literatures broadly focus on persuasion problems and delegation problems, respectively. Kolotilin and Zapechelnyuk (2018) explore the mathematical connections

[1] A closely related literature studies incentive schemes for out-sourced research and innovation processes (see, e.g., Benkert and Letina, 2020; Bimpikis et al., 2019; Halac et al., 2016; Manso, 2011).

between such problems. Boleslavsky and Kim (2018) and Göx and Michaeli (2019) study persuasion by a principal, whereas I study persuasion by an agent.

2 The contracting game

This section describes a contracting game played by a recruitment agent and their client, who may be either a worker or a firm. In each period of the game, the client pays the agent if they find the client a match. Search costs depend on market thickness, which the agent observes but the client does not observe. The agent and client are risk-neutral, and choose search efforts and payments to maximize their expected payoffs. This maximization delivers a belief-contingent optimal contract.

2.1 Game description

The labor market is either thick ($\omega = 0$) or thin ($\omega = 1$). The client and agent have common prior $\Pr(\omega = 1) = \pi_1 \in (0, 1)$, which also defines the true distribution of ω. The client does not observe ω. The agent observes ω and chooses how much effort $e \geq 0$ to allocate to finding a match. If the market is thick then search succeeds if and only if $e \geq e_L > 0$. If the market is thin then search succeeds if and only if $e \geq e_H > e_L$. The effort cost function $c(e)$ is increasing in e, with $c(0) = 0$, $c_L \equiv c(e_L) \in (0, 1)$, and $c_H \equiv c(e_H) \in (c_L, 1)$.

The client knows c_L and c_H. They observe whether search succeeds but not the agent's effort. The client pays the agent $p_t \in (0, 1]$ if search succeeds in period t. If search fails then the client pays nothing and the game continues for another period. The client and agent have inter-temporal discount factor $\delta \in (0, 1)$. I assume the market for recruitment agents is sufficiently competitive that the agent cannot profit from choosing $e = 0$ when the client's payment p_t covers the marginal effort cost of search. Therefore, the agent's period t payoff is

$$\psi_t(e, \omega) \equiv \begin{cases} -c(e) & \text{if } e < e_L \\ p_t - c(e) & \text{if } e_L \leq e < e_H \text{ and } \omega = 0 \\ -c(e) & \text{if } e_L \leq e < e_H \text{ and } \omega = 1 \\ p_t - c(e) & \text{if } e \geq e_H. \end{cases}$$

Given p_t, the agent maximizes $\psi_t(e, \omega)$ by choosing $e \in \{0, e_L\}$ when the market is thick ($\omega = 0$) and $e \in \{0, e_H\}$ when it is thin ($\omega = 1$). The client anticipates these choices of e when choosing p_t.

The client's period t payoff is

$$\phi_t(p_t,\omega) \equiv \begin{cases} 0 & \text{if } p_t < c_L \\ Q - p_t & \text{if } c_L \le p_t < c_H \text{ and } \omega = 0 \\ 0 & \text{if } c_L \le p_t < c_H \text{ and } \omega = 1 \\ q - p_t & \text{if } p_t \ge c_H, \end{cases}$$

where $Q \in (c_H, 1]$ and $q \in (c_H, Q]$ are the mean qualities of a match when $\omega = 0$ and $\omega = 1$. The client does not observe ω, but realizes that paying $p_1 = c_L$ strictly dominates paying $p_1 \in (c_L, c_H)$ or $p_1 < c_L$, and that paying $p_1 = c_H$ strictly dominates paying $p_1 > c_H$. Thus, there are two cases to consider:

1. Suppose $p_1 = c_H$. The agent finds a match independently of ω. The client receives expected payoff

   $$\Phi(c_H) \equiv (1 - \pi_1)Q + \pi_1 q - c_H, \tag{1}$$

   the agent receives expected payoff

   $$\Psi(c_H) \equiv (1 - \pi_1)(c_H - c_L), \tag{2}$$

   and the game ends.

2. Suppose $p_1 = c_L$. If $\omega = 0$ then the agent finds a match, the client receives payoff $(Q - c_L)$, the agent receives a payoff of zero, and the game ends. If $\omega = 1$ then the agent does not find a match and the game continues for a second period. This continuation reveals to the client that $\omega = 1$ (for otherwise paying $p_1 = c_L$ would have yielded a match), so they revise their belief to $\pi_2 = 1$. Accordingly, the client pays $p_2 = c_H$ and the agent finds a match. The client receives period payoff $(q - c_H)$, the agent receives a period payoff of zero, and the game ends. Therefore, the client's discounted expected payoff from choosing $p_1 = c_L$ is

   $$\Phi(c_L) \equiv (1 - \pi_1)(Q - c_L) + \pi_1 \delta (q - c_H), \tag{3}$$

   and the agent's expected payoff from the client choosing $p_1 = c_L$ is

   $$\Psi(c_L) \equiv 0. \tag{4}$$

Assuming the market for recruitment agents is competitive prevents the agent from rejecting the low payment $p_1 = c_L$ when $\omega = 0$ in hope that the client will offer $p_t = c_H$ in some future period t, which would increase the agent's payoff. I relax this competition assumption in Section 4.

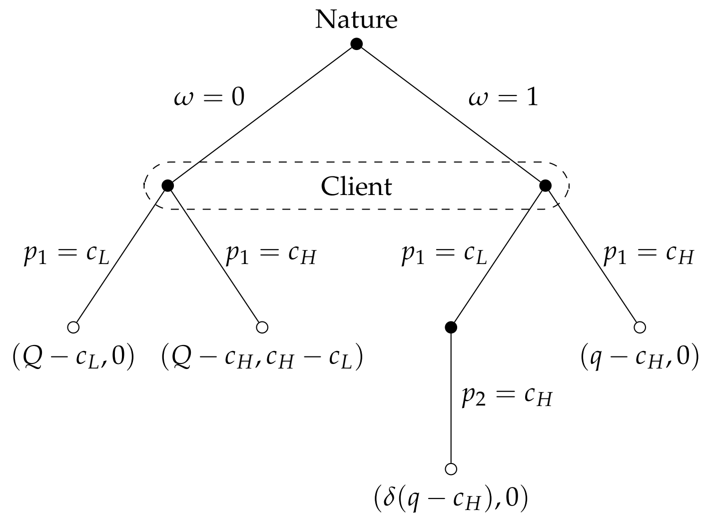

Figure 1: Extensive form of contracting game after deleting dominated strategies. Dash-bordered region represents client's information set when choosing p_1. State-contingent payoffs in brackets (client's first).

Figure 1 summarizes the extensive form of the contracting game after deleting strictly dominated strategies. Nature chooses a state $\omega \in \{0, 1\}$ that determines the effort $e \in \{e_L, e_H\}$ needed to find a match. Without observing ω, the client chooses a first-period payment $p_1 \in \{C_L, C_H\}$. If $(p_1, \omega) \neq (c_L, 1)$ then the agent finds a match, the agent and client receive their payoffs, and the game ends. If $(p_1, \omega) = (c_L, 1)$ then the game proceeds for another period. The client infers that $\omega = 1$ and chooses $p_2 = c_H$. The agent finds a match, the agent and client receive their payoffs, and the game ends.

2.2 The optimal contract

The client maximizes their *ex ante* expected payoff by choosing

$$p_1^* \in \underset{p_1 \in \{c_L, c_H\}}{\arg\max} \; \Phi(p_1),$$

breaking ties in favor of choosing $p_1 = c_H$.[2] It follows from (1) and (3) that

$$p_1^* = \begin{cases} c_L & \text{if } \pi_1 < \pi^* \\ c_H & \text{otherwise,} \end{cases}$$

[2] I assume this tie-breaking rule because choosing $p_1 = c_H$ delivers a certain payoff but choosing $p_1 = c_L$ delivers a random payoff, so if both choices have the same expected payoff then a risk-averse client would prefer $p_1 = c_H$.

where

$$\pi^* \equiv \frac{c_H - c_L}{\delta c_H - c_L + (1-\delta)q} \tag{5}$$

is the threshold belief at which the client is indifferent. Intuitively, the worker prefers to gamble that low effort is sufficient to find a match whenever they are sufficiently confident the market is thick and effort costs are low.

The optimal contract is belief-contingent: if $\pi_1 < \pi^*$ then it stipulates a first-period payment c_L for effort e_L if the market is thick and a second-period payment c_H for effort e_H if the market is thin; if $\pi_1 \geq \pi^*$ then it stipulates a first-period payment c_H for effort e_H. This belief-contingency comes from the client being less willing to pay for the agent's services when the market is thick than when the market is thin.

The threshold belief π^* is increasing in c_H and decreasing in c_L. Intuitively, the higher is the relative cost of searching in a thin market, the more pessimistic about market thickness the client must be to forgo gambling that effort costs are low. Likewise, π^* is decreasing in δ and q, since if the client is more patient or thin-market match prospects are worse then the client is more willing to delay matching in the thin market until the second period. The threshold π^* is independent of Q because the client faces the same *ex ante* probability of receiving a match of quality Q regardless of which payment $p_1 \in \{c_L, c_H\}$ they choose.

3 Persuasion opportunities

The agent receives expected payoff $\Psi(c_H) > 0$ when the client chooses $p_1 = c_H$ and certain payoff $\Psi(c_L) = 0$ when the client chooses $p_1 = c_L$. Thus, the agent strictly prefers $p_1 = c_H$ to $p_1 = c_L$. The client shares this preference when $\pi_1 > \pi^*$. But if $\pi_1 < \pi^*$ then the client prefers $p_1 = c_L$. The agent can manipulate this preference using Bayesian persuasion (Kamenica and Gentzkow, 2011). They can send signals of market conditions that induce the client to choose $p_1 = c_H$ with non-zero probability, making the agent strictly better off than without such signaling.

This section analyzes how the agent can benefit from manipulating their client's beliefs about market thickness, captured by π, and prospective match qualities, captured by q. I conduct these analyses separately in the next two subsections.

3.1 Manipulating beliefs about market thickness

Let $\pi_1 < \pi^*$. The agent sends a signal $s \in \{0, 1\}$ with state-contingent distribution

$$\Pr(s \mid \omega) = \begin{cases} 1 - x & \text{if } (s, \omega) = (0, 0) \\ x & \text{if } (s, \omega) = (1, 0) \\ s & \text{if } \omega = 1 \end{cases} \tag{6}$$

for some $x \in [0, 1]$. Upon observing s, the client uses Bayes' formula to form a posterior belief

$$\begin{aligned} \tilde{\pi}_1(s) &\equiv \Pr(\omega = 1 \mid s) \\ &= \frac{\Pr(s \mid \omega = 1) \Pr(\omega = 1)}{\Pr(s)} \\ &= \begin{cases} 0 & \text{if } s = 0 \\ \frac{\pi_1}{\pi_1 + x(1 - \pi_1)} & \text{if } s = 1 \end{cases} \end{aligned}$$

about the distribution of ω. If the client observes $s = 0$ then they choose $p_1 = c_L$ since $\tilde{\pi}_1(0) < \pi^*$. If they observe $s = 1$ then they choose $p_1 = c_H$ if $\tilde{\pi}_1(1) \geq \pi^*$ and $p_1 = c_L$ otherwise. Therefore, the agent's expected payoff from providing the signal before playing the contracting game is

$$\Pi(x) \equiv \begin{cases} \Psi(c_H) \Pr(s = 1) & \text{if } \tilde{\pi}_1(1) \geq \pi^* \\ 0 & \text{otherwise.} \end{cases}$$

Since $\Pr(s = 1) = \pi_1 + x(1 - \pi_1)$ increases with x but $\tilde{\pi}_1(1)$ decreases with x, the maximizer x^* of $\Pi(x)$ satisfies $\tilde{\pi}_1(1) = \pi^*$, which implies

$$x^* = \frac{\pi_1(1 - \pi^*)}{\pi^*(1 - \pi_1)}.$$

Setting $x = x^*$ gives the agent expected payoff

$$\Pi(x^*) = \frac{\pi_1(1 - \pi_1)(c_H - c_L)}{\pi^*},$$

which is strictly positive and thus exceeds their payoff without the signal. Intuitively, observing $s = 1$ increases the client's willingness to pay for the agent's services because it makes the client believe the market is thin and search costs are high. The client forms these beliefs despite knowing the signal structure and that it was designed to be maximally persuasive.

The expected payoff $\Pi(x^*)$ is decreasing in π^*, which is decreasing in q. Thus, holding beliefs π_1 constant, persuading the client to revise their beliefs about market thickness is more beneficial (in expectation) when match prospects in a thin market are better. Intuitively, the better are those prospects, the less willing is the client to delay a thin-market match and the more scope the agent has to exploit this delay aversion.

3.2 Manipulating beliefs about match qualities

Let $\pi_1 < \pi^*$. For simplicity, suppose q equals either $q_H \in (c_H, 1]$ or $q_L \in (c_H, q_H)$. The agent and client have common prior $\Pr(q = q_H) = \tau \in (0,1)$. The agent observes q but the client does not. The client sets π^* under the assumption that $q = \tau q_H + (1-\tau)q_L$. Since π^* falls with q, the agent wants to shift τ upward until $\pi_1 \geq \pi^*$, in which case the client prefers to pay $p_1 = c_H$. The agent achieves this goal by sending a signal $s' \in \{0,1\}$ with state-contingent distribution

$$\Pr(s' \mid q) = \begin{cases} 1-y & \text{if } (s', q) = (0, q_L) \\ y & \text{if } (s', q) = (1, q_L) \\ s' & \text{if } q = q_H \end{cases} \tag{7}$$

for some $y \in [0,1]$. Upon observing s', the client uses Bayes' formula to form a posterior belief

$$\begin{aligned} \tilde{\tau}(s') &\equiv \Pr(q = q_H \mid s') \\ &= \begin{cases} 0 & \text{if } s' = 0 \\ \frac{\tau}{\tau + y(1-\tau)} & \text{if } s' = 1 \end{cases} \end{aligned} \tag{8}$$

about the distribution of q. This belief, in turn, induces posterior estimates

$$\tilde{q}(s') \equiv \tilde{\tau}(s')q_H + (1 - \tilde{\tau}(s'))q_L \tag{9}$$

of the mean match quality q and

$$\tilde{\pi}^*(s') \equiv \frac{c_H - c_L}{\delta c_H - c_L + (1-\delta)\tilde{q}(s')} \tag{10}$$

of the threshold belief π^*. Now $\tilde{\pi}^*(s')$ is decreasing in $\tilde{q}(s')$, but $\tilde{q}(s')$ is increasing $\tilde{\tau}(s')$, so $\tilde{\pi}^*(s')$ is also decreasing $\tilde{\tau}(s')$. If the client observes $s' = 0$ then they choose $p_1 = c_L$ because $\tilde{\tau}(0) = 0 < \tau$ and hence $\tilde{\pi}^*(0) > \pi^* > \pi_1$. If the client observes $s' = 1$ then they choose $p_1 = c_H$ if $\pi_1 \geq \tilde{\pi}^*(1)$ and $p_1 = c_L$ otherwise. Thus, the agent's expected payoff from providing the signal before playing the contracting game is

$$T(y) \equiv \begin{cases} \Psi(c_H)\Pr(s' = 1) & \text{if } \pi_1 \geq \tilde{\pi}^*(1) \\ 0 & \text{otherwise.} \end{cases}$$

Since $\Pr(s' = 1) = \tau + y(1-\tau)$ and $\tilde{\pi}^*(1)$ both increase with y, the maximizer y^* of $T(y)$ satisfies $\pi_1 = \tilde{\pi}^*(1)$. It follows from (8)–(10) that

$$y^* = \frac{\tau(q_H - \tilde{q}(1))}{(1-\tau)(\tilde{q}(1) - q_L)},$$

where

$$\tilde{q}(1) = \frac{c_H(1-\pi_1\delta) - c_L(1-\pi_1)}{\pi_1(1-\delta)}$$

is the client's posterior estimate of q after observing $s' = 1$ when $y = y^*$. Then $\tilde{q}(1) > q_L = \tilde{q}(0)$. Setting $y = y^*$ gives the agent expected payoff

$$T(y^*) = \frac{\tau(q_H - q_L)(1-\pi_1)(c_H - c_L)}{\tilde{q}(1) - q_L},$$

which is strictly positive and so exceeds their payoff without the signal. Thus, the agent benefits from manipulating the client's beliefs about q. This benefit is increasing in the client's belief π_1 that the market is thin. Intuitively, the larger is π_1, the less attractive the thin market needs to be for the client to prefer paying $p_1 = c_H$, and, hence, the less constrained is the signal the agent must send to make the high payment sufficiently rewarding in expectation.

4 Relaxing the competition assumption

The contracting game presented in Section 2 assumes the agent cannot charge markups for their services due to competitive pressure from other agents. This assumption prevents the game from extending beyond two periods because the client infers the state perfectly upon reaching the second period and chooses $p_2 = c_H$, ending the game.

Suppose the agent can choose $e = 0$ in response to payments less than c_H. Such a choice gives the agent and client period payoffs of zero. Therefore, if the agent can choose $e = 0$ then the client can receive a positive payoff only by choosing $p_T = c_H$ in some period T. The *ex ante* present value of this payoff is $\delta^{T-1}((1-\pi_1)Q + \pi_1 q - c_H)$, which is maximized when $T = 1$. Thus, the client's best response to the agent's ability to reject low payments is to choose $p_1 = c_H$ and end the game after a single period. I consider this scenario *very* degenerate: if the client always chooses $p_1 = c_H$ then they face no interesting trade-offs and the agent has no incentive to persuade.

Thus, for the contracting game to be interesting, there need to be forces constraining the agent's ability to charge markups. Competition is one such force. This section discusses another force: the ability for the client to find matches on their own.

4.1 Introducing self-search as an outside option

Suppose the agent can choose $e = 0$ when the market is thick and $p_t = c_L$. They make this choice with probability $\theta \in [0, 1]$—the agent's "shirk rate"—independently for each search period.[3] The client constrains the agent by threatening to find matches on their own. However, search effort is

[3] One could make θ time- and/or belief-dependent. I assume independence for simplicity and tractibility.

more costly for the client: they must pay $\kappa \in (c_L, c_H)$ in self-search costs to find a match when the market is thick, and some amount larger than c_H when the market is thin. Thus, the client prefers to hire the agent when the market is thin, but prefers to self-search if the market is thick and the agent shirks sufficiently often.

Let π_t denote the client's period t belief that $\omega = 1$, and let $V_\theta(\pi_t)$ denote the present value of the client's expected payoff from holding that belief when the shirk rate is θ and the client chooses payoff-maximizing payments. The client has three period t choices:

1. Pay $p_t = c_H$ to guarantee a match and receive expected payoff $(1-\pi_t)Q + \pi_t q - c_H$.
2. Self-search at effort cost κ, which yields a match of mean quality Q with probability $(1-\pi_t)$. If no match is found then the client infers $\omega = 1$ (for otherwise κ would have been sufficient to find a match), so in period $(t+1)$ they pay the agent $p_{t+1} = c_H$ for finding a match. This strategy gives the client expected payoff $(1-\pi_t)(Q-\kappa) + \pi_t\delta(q - c_H)$.
3. Pay $p_t = c_H$, which yields a match of mean quality Q if the market is thick and the agent does not shirk. Otherwise, no match is found, the client updates their belief to
$$\pi_{t+1} = \frac{\pi_t}{(1-\pi_t)\theta + \pi_t},$$
and the game continues for another period. Thus, the client's expected payoff from choosing $p_t = c_L$ is $(1-\pi_t)(1-\theta)(Q-c_L) + ((1-\pi_t)\theta + \pi_t)\delta V_\theta(\pi_{t+1})$.

Given these three choices, the present value $V_\theta(\pi_t)$ must satisfy the Bellman equation

$$V_\theta(\pi_t) = \max\left\{\begin{array}{l} (1-\pi_t)Q + \pi_t q - c_H, \\ (1-\pi_t)(Q-\kappa) + \pi_t\delta(q-c_H), \\ (1-\pi_t)(1-\theta)(Q-c_L) + ((1-\pi_t)\theta+\pi_t)\delta V_\theta\left(\frac{\pi_t}{(1-\pi_t)\theta+\pi_t}\right) \end{array}\right\}. \tag{11}$$

Notice that $V_\theta(1) = q - c_H$ for all $\theta \in [0,1]$, since if the client knows the market is thin then they cannot do better than pay $p_t = c_H$ to guarantee a match.

4.2 Numerical patterns

The client's optimal period t strategy depends only on their belief π_t about ω. Consequently, the client's choices follow a Markov decision process and the Bellman equation (11) can be solved using value function iteration (Bellman, 1957). I use this algorithm to approximate $V_\theta(\pi_t)$ for beliefs $\pi_t \in \{0, 0.001, 0.002, \ldots, 1\}$, fixing the parameters $(Q, q, c_L, c_H, \delta, \kappa) = (1, 0.8, 0.5, 0.6, 0.95, 0.55)$ but varying the agent's shirk rate θ.

Figure 2 plots $V_\theta(\pi_t)$ for each $\theta \in \{0, 0.2, 0.4, 0.6, 0.8, 1\}$. The value $V_\theta(\pi_t)$ falls with π_t because

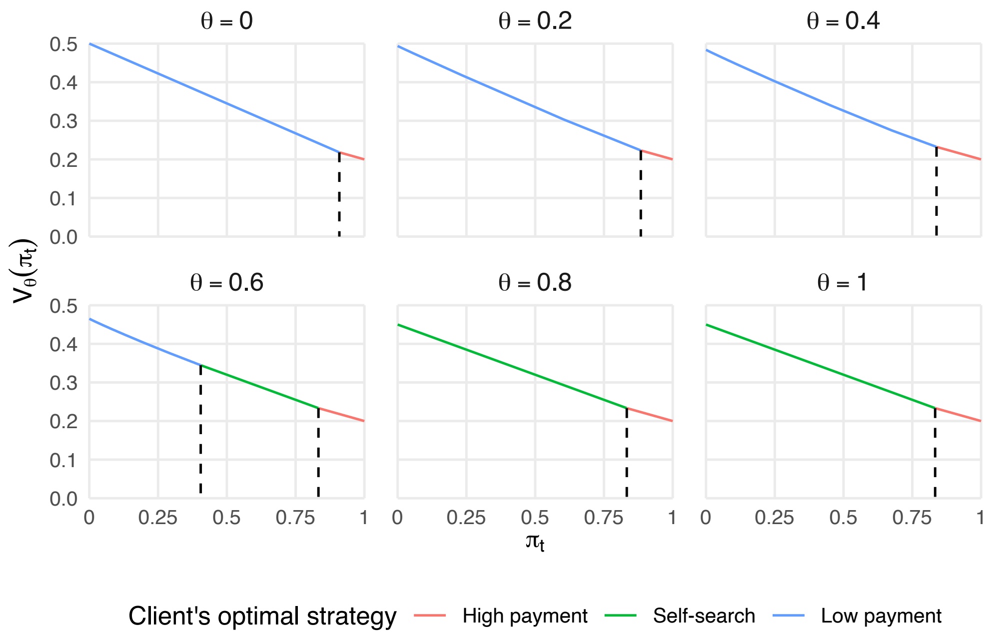

Figure 2: Values $V_s(\pi_t)$ for $(Q, q, c_L, c_H, \delta, \kappa) = (1, 0.8, 0.5, 0.6, 0.95, 0.55)$ and varying θ. Dashed lines indicate threshold beliefs at which client is indifferent between strategies.

if the client is more certain the market is thin then they are more likely to make the high payment $p_t = c_H$ and get a lower quality match. If the agent never shirks ($\theta = 0$) then the client uses the same strategy as in Section 2: pay $p_t = c_L$ if $\pi_t < \pi^*$ and $p_t = c_H$ otherwise. On the other hand, if the agent *always* shirks ($\theta = 1$) then the client chooses between self-search and paying $p_t = c_H$ to guarantee a match. The client prefers self-search if and only if $\pi_t < \pi^\dagger$, where

$$\pi^\dagger \equiv \frac{c_H - \kappa}{\delta c_H - \kappa + (1 - \delta)q}$$

is the threshold belief at which the client is indifferent between self-search and paying $p_t = c_H$.

The threshold belief π^θ at which the client is indifferent between paying $p_t = c_L$ and $p_t = c_H$ falls as the agent's shirk rate θ rises. Intuitively, if the agent is more likely to shirk then the client has less bargaining power, making them more willing to "give in" and pay $p_t = c_H$. The threshold π^θ coincides with $\pi^\dagger$ at some shirk rate $\theta^\dagger$ before falling to zero at some shirk rate $\theta^\ddagger$. The client chooses between paying $p_t = c_L$ and $p_t = c_H$ when $\theta < \theta^\dagger$, and between self-search and paying $p_t = c_H$ when $\theta > \theta^\ddagger$. They choose among all three strategies when $\theta \in [\theta^\dagger, \theta^\ddagger]$. All choices depend on the client's belief π_t, as indicated by lines' colors in Figure 2.

Now consider the agent's "optimal shirking" problem: choosing the shirk rate θ to maximize their expected payoff in the (perfect) Bayesian Nash equilibrium of the modified contracting game resulting from their choice. When solving this problem, the agent trades off setting θ low enough to prevent the client from self-searching (which always gives the agent zero profit) and setting θ high enough to avoid accepting low payments when the market is thick. I defer optimizing this trade-off to future work.

4.3 Persuasion benefits

If the agent shirks with probability θ and the client can self-search, then the client's optimal strategy is to pay $p_t = c_H$ whenever their belief π_t that the market is thin exceeds $\max\{\pi^\theta, \pi^\dagger\}$. Given such beliefs, the agent cannot benefit from persuasion because the client's default action already gives the agent the highest possible expected payoff. On the other hand, if $\pi_t < \max\{\pi^\theta, \pi^\dagger\}$ then the client's default action is *not* to pay $p_t = c_H$, and so the agent can benefit in expectation from manipulating the client's beliefs about market thickness and prospective match qualities.

Suppose $\pi_t < \max\{\pi^\theta, \pi^\dagger\}$. Now $\pi^\theta \leq \pi^*$ with equality if and only if $\theta = 0$, whereas $\pi^\theta \geq \pi^\dagger$ for $\theta \leq \theta^\dagger$ and $\pi^\theta < \pi^\dagger$ for $\theta > \theta^\dagger$. But $\pi^\dagger < \pi^*$ as $\kappa > c_L$. Thus $\max\{\pi^\theta, \pi^\dagger\} \leq \pi^*$ with equality if and only if $\theta = 0$. Consequently, if the agent sends a signal $s \in \{0, 1\}$ of market thickness with distribution (6), then the maximal expected payoff

$$\frac{\pi_1(1-\pi_1)(c_H - c_L)}{\max\{\pi^\theta, \pi^\dagger\}}$$

is weakly greater than the maximal expected payoff $\Pi(x^*)$ attainable when the the agent cannot shirk (and strictly greater when $\theta > 0$). Likewise, if the agent sends a signal $s' \in \{0, 1\}$ of mean match quality with distribution (7), then the maximal expected payoff is (weakly) greater than the maximal expected payoff $T(y^*)$ attainable when the the agent cannot shirk. In this way, being able to shirk makes persuasion more beneficial for the agent. This gain in persuasion benefits comes from the agent exploiting their search cost advantage over the client when the market is thick.

5 Extensions

Section 4 extended the contracting game described in Section 2 by relaxing the competitive forces on the agent and allowing the client to find matches on their own. This section ends the paper by describing three more extensions to be considered for future study.

First, one could let the mean match qualities Q and q depend on the agent's search effort, and the contracted payment depend on the realized match quality. Then the agent would trade-off finding a "bad match" quickly with a "good match" slowly (as discussed, e.g., by Zorc et al., 2019), and the client would choose a payment structure to elicit their optimal trade-off. This extension hints at the agent's role as a matching mechanism, entering contracts with one side of the market that condition on opportunities available on the other side.

Second, one could make the agent's "matching mechanism" role explicit: let them enter contracts with workers and firms simultaneously. This would connect the analysis to the literatures on two-sided matching and platform design. It would also allow one to explore the agent's role as an informational intermediary: workers and firms could "advertise" to each other by sending signals of their quality through the agent.

Third, one could consider the agent's incentives to build a reputation *across* clients. Reputation concerns offer another persuasion motive: the agent wants clients to believe the market is thinner than reality so that, upon failing to find a match, the agent has plausible deniability (and preserves their reputation) if they justify their failure by appealing to market thinness. One could also allow for heterogeneity in reputational priorities, opening the door to other inefficiencies. For example, prioritizing match quality may lead to sub-optimally long searches, while prioritizing turnover may lead to sub-optimally low quality matches.

References

Arnold, M. A. (1992). The Principal-Agent Relationship in Real Estate Brokerage Services. *Real Estate Economics*, 20(1):89–106.

Artemov, G. (2021). Assignment mechanisms: Common preferences and information acquisition. *Journal of Economic Theory*, 198:105370.

Ashlagi, I., Braverman, M., Kanoria, Y., and Shi, P. (2020). Clearing Matching Markets Efficiently: Informative Signals and Match Recommendations. *Management Science*, 66(5):2163–2193.

Bellman, R. (1957). A Markovian Decision Process. *Journal of Mathematics and Mechanics*, 6(5):679–684.

Benkert, J.-M. and Letina, I. (2020). Designing Dynamic Research Contests. *American Economic Journal: Microeconomics*, 12(4):270–289.

Bikhchandani, S. (2017). Stability with one-sided incomplete information. *Journal of Economic Theory*, 168:372–399.

Bimpikis, K., Ehsani, S., and Mostagir, M. (2019). Designing Dynamic Contests. *Operations Research*, 67(2):339–356.

Boleslavsky, R. and Kim, K. (2018). Bayesian Persuasion and Moral Hazard.

Chen, Y.-C. and Hu, G. (2020). Learning by matching. *Theoretical Economics*, 15(1):29–56.

Göx, R. F. and Michaeli, B. (2019). Optimal Information Design and Incentive Contracts with Performance Measure Manipulation.

Halac, M., Kartik, N., and Liu, Q. (2016). Optimal Contracts for Experimentation. *Review of Economic Studies*, 83(3):1040–1091.

Kamenica, E. and Gentzkow, M. (2011). Bayesian Persuasion. *American Economic Review*, 101(6):2590–2615.

Kanoria, Y. and Saban, D. (2021). Facilitating the Search for Partners on Matching Platforms. *Management Science*, 67(10):5990–6029.

Kolotilin, A. and Zapechelnyuk, A. (2018). Persuasion Meets Delegation. Discussion Paper 2018-06, School of Economics, University of New South Wales.

Levitt, S. D. and Syverson, C. (2008). Market Distortions When Agents Are Better Informed: The Value of Information in Real Estate Transactions. *Review of Economics and Statistics*, 90(4):599–611.

Lewis, T. R. (2012). A theory of delegated search for the best alternative. *RAND Journal of Economics*, 43(3):391–416.

Lewis, T. R. and Ottaviani, M. (2008). Search Agency.

Liu, Q., Mailath, G. J., Postlewaite, A., and Samuelson, L. (2014). Stable Matching With Incomplete Information. *Econometrica*, 82(2):541–587.

Manso, G. (2011). Motivating Innovation. *Journal of Finance*, 66(5):1823–1860.

Postl, P. (2004). Delegated Search: Procedure Matters. Discussion Paper 04-17, Department of Economics, University of Birmingham.

Ulbricht, R. (2016). Optimal delegated search with adverse selection and moral hazard. *Theoretical Economics*, 11(1):253–278.

Zorc, S., Tsetlin, I., and Hasija, S. (2019). The When and How of Delegated Search. Darden Business School Working Paper 3463626, University of Virginia.