ChatGPT in Directed Search

LI, HAO
University of British Columbia

February 18, 2026*

Abstract. Job market papers can sometimes be made appear well-written by the use of ChatGPT regardless of actual quality of candidates. This makes a job market paper a more accurate signal when the quality of the job candidate who authored the paper is high, but a less accurate signal when the candidate's quality is low, and generally has an ambiguous effect on the average informativeness of the job market paper depending on the prior belief about the candidate quality. To illustrate the impact of ChatGPT in a search market, we use the a simplified version of Peters (2010), where departments with private information about the common value of their own job offers to job candidates make targeted offers depending on whether candidates' job market papers are well-written or poorly-written. We show that when ChatGPT is informationally ambiguous, it makes all departments worse off, and can at the same time make all candidates better off.

*I would like to thank Mike Peters for helpful discussions.

1 Introduction

In the job market for economists, job candidates can use ChatGPT to make their job market papers appear well-written regardless of the actual quality of the job candidates who wrote the papers. ChatGPT increases the number of job market papers that are well-written, but has two opposing information effects: it simultaneously makes the job market paper a more accurate signal when the underlying quality of the job candidates is high, and it makes the job market paper a less accurate signal when the underlying quality is low. What is impact of the use of ChatGPT on the search market? What is the impact on welfare of departments and candidates? We argue that when the effect of ChatGPT on the informativeness of the job market paper is on average "ambiguous" because it depends on the prior belief about the quality of the job candidate, the use of ChatGPT makes all departments worse off, and can at the same time make all candidates better off.

We show that when the use of ChatGPT increases the likelihood ratio of a high-quality candidate with a well-written paper relative to a low-quality candidate, it improves the average informativeness of job market papers regardless of the prior belief about the candidate's quality as measured by conditional entropy. In this case of "informative" ChatGPT, it improves the perceived quality of a candidate with a well-written paper and reduces the perceived quality of one with a poorly-written paper. Conversely, when ChatGPT increases the likelihood ratio of a high-quality candidate with a poorly-written paper relative to a low-quality candidate, it reduces the average informativeness of job market papers regardless of the prior belief about the candidate's quality. This is the case of "uninformative ChatGPT," where it reduces the perceived quality of a candidate with a well-written paper and improves the perceived quality of one with a poorly-written paper. In between the two cases, ChatGPT is informationally "ambiguous," as its effect on the average informativeness of job market papers depends on the prior belief about the candidate's quality. In this case, ChatGPT reduces the perceived quality of a candidate both with a well-written paper and a poorly-written paper.

We use a simplified version of directed search model of Peters (2010) to evaluate effects of ChatGPT. In this model, departments have private information about the common value of their own offers to the other side, and this value is independently drawn from a single continuous distribution. They make targeted offers to candidates whose job market papers are either well-written or poorly-written. There is a unique symmetric equilibrium with a threshold offer value where all departments with offers higher than the threshold target candidates with well-written papers exclusively, and those with offers below the threshold target the same candidates with a common probability that is independent of their offer values.

The impact of ChatGPT on the unique symmetric equilibrium depends only on its effect on the ratio of the perceived quality of a candidate with a well-written paper to that of a poorly-written paper. A greater ratio reduces the equilibrium threshold offer value as it becomes relatively more attractive for departments with high offers to target candidates with a well-written paper. At the same time, a greater ratio reduces the equilibrium probability that departments with low offers target candidates with a well-written paper, because they face greater competition for those candidates. These two effects combined, more high offers and fewer low offers go to candidates with a well-written paper, and more low offers go to candidates with a poorly-written paper. Thus, a greater ratio of the perceived quality of the candidate with a well-written paper to that of a poorly-written paper shifts the equilibrium in favor of candidates with a well-written paper.

Since informative ChatGPT improves the perceived quality of a candidate with a well-written paper and reduces the perceived quality of a candidate with a poorly-written paper, it shifts the equilibrium in favor of candidates with a well-written paper. In contrast, since it has exactly opposite effects on the perceived qualities of candidates with a well-written paper and a poorly-written paper, uninformative ChatGPT shifts the equilibrium in favor of candidates with a poorly-written paper. In the case of ambiguous ChatGPT, the perceived qualities of candidates with a well-written paper and a poorly-written paper are both reduced, and the ratio of the perceived qualities can go either way. The ratio is more likely to increase, and thus its impact on the equilibrium is more likely to favor candidates with a well-written paper, if ChatGPT on average increases the informativeness of the job market paper signal.

The impact of ChatGPT on equilibrium welfare of candidates is entirely determined by whether it shifts the equilibrium in favor of those with a well-written paper or with a poorly-written paper. The effects on candidates with a well-written paper and those with a poorly-written paper are always exactly opposite. Thus, ChatGPT is more likely to make candidates with a well-written paper better off and make those with a poorly-written paper worse off, if it on average increases the informativeness of the job market paper signal. From the ex ante perspective, since ChatGPT always generates more well-written papers regardless of the underlying quality, and since candidates with a well-written paper are better off relative to those with a poorly-written paper, the ex ante impact ChatGPT is to make candidates better off. Thus, when ChatGPT is ambiguous and causes little shift in the offer distribution in equilibrium, all candidates are better off ex ante.

The impact of ChatGPT on equilibrium welfare of an individual department depends on whether it has an offer higher than the equilibrium threshold or one lower than the threshold. Since a department with an offer higher than the equilibrium threshold targets only candidates with a well-rewritten paper, it is better off when and only when ChatGPT improves the perceived quality of a candidate with a well-written paper. As a result, it is better off with informative ChatGPT, but worse off with either ambiguous ChatGPT or uninformative ChatGPT. Since departments with offers below the threshold mix between targeting candidates with a well-written paper and those with a poorly written paper, the impact of ChatGPT generally depends on the comparison of its effects on the perceived qualities of a candidate with a well-written paper and poorly-written paper. From the ex ante perspective, the effect of ChatGPT on the welfare of departments with offer below the threshold also depends on the realized numbers of candidates with a well-written paper versus a poorly-written paper. However, when ChatGPT is ambiguous, regardless of the realized numbers, these departments, as well as departments with offers above the threshold, are unambiguously worse off because ChatGPT reduces both perceived qualities.

2 A Model of Directed Search

There are $D \geq 2$ economics departments and $C \geq 2$ job candidates. Each economic department d, $d = 1, \dots, D$, draws an offer x_d from a distribution $F(\cdot)$ with support $[0, \overline{x}]$, where x_d is for the payoff to a candidate who accepts the offer from d. We assume that x_d is independent across D departments.

Each job candidate $c = 1, \dots, C$, has an unobservable value V_c, which is the payoff to each department if the candidate accepts an offer from the department. We assume that the random variable V_c is identical, equal to V, and independent across candidates. For simplicity, assume that V is binary, with support $v_H > v_L > 0$. We refer to v_H as "high quality" and v_L as "low quality." For each $i = H, L$, let the common prior probability that $V_c = v_i$ be $p_i > 0$, with $p_H + p_L = 1$. The publicly observed signal σ about V_c is also assumed to be a binary random variable, which models the job market paper of candidate c. We denote the two signals as $\sigma = \sigma^h, \sigma^l$. Let

$$q_H^h = \Pr(\sigma^h | v_H),$$

and

$$q_L^h = \Pr(\sigma^h | v_L).$$

We assume that

$$q_H^h > q_L^h,$$

so that σ^h is the "high" signal, or a well-written paper, and σ^l the "low" signal, or a poorly-written paper. The greater is q_H^h, the more accurate is signal σ^h; the smaller is q_L^h for v_H, the more accurate is signal σ^l for v_L. Define

$$\pi^h = p_H q_H^h + p_L q_L^h$$

as the probability that a given candidate is observed to have signal σ^h, and define $\pi^l = 1 - \pi^h$ as the probability that the signal is σ^l. The probability of $V_c = v_H$ conditional on σ^h as β_H^h, given by

$$\beta_H^h = \frac{p_H q_H^h}{\pi^h}$$

and analogously define the probability $V_c = v_H$ conditional on σ^l as

$$\beta_H^l = \frac{p_H(1 - q_H^h)}{\pi^l}.$$

Since $q_H^h > q_L^h$, we have $\beta_H^h > \beta_H^l$. The posterior mean value given signal σ_j, $j = h, l$, is

$$u^j = \beta_H^j v_H + (1 - \beta_H^j) v_L = v_H - (1 - \beta_H^j)(v_H - v_L).$$

Since $\beta_H^h > \beta_H^l$, we have

$$u^h > u^l.$$

By law of iterated expectations,

$$\pi^h u^h + \pi^l u^l = p_H v_H + p_L v_L.$$

The game proceeds as follows.

- Each candidate c, $c = 1, \ldots, C$, independently draws a public signal, σ^h or σ^l, from from V_c.
- Each department d, $d = 1, \ldots, D$, independently draws a private offer x_d from distribution $F(\cdot)$.

- All departments each simultaneously choose a candidate to make an offer to.
- For each candidate c who accepts offer x_d among the offers they have received, c's payoff is x_d and d's payoff is the posterior mean of V_c conditional on c's signal. A candidate who has received no offers, or a candidate who rejects all offers they have received, gets a payoff of 0; a department d whose offer x_d is rejected by the candidate c that d made the offer to gets a payoff of 0.

2.1 ChatGPT

Let $\mathrm{d}q_i^h$ be the infinitesimal change in q_i^h in each $i = H, L$. To model ChatGPT, consider

$$\mathrm{d}q_i^h > 0$$

for each $i = H, L$. We have

$$\mathrm{d}\pi^h > 0,$$

i.e., ChatGPT generates more candidates with signal σ^h.

ChatGPT generally has an ambiguous impact on the informativeness of public information about job market candidates. It has two opposing effects on signal accuracy: it simultaneously increases q_H^h and q_L^h. When the candidate is actually high quality, ChatGPT helps by generating signal σ^h with a greater probability; when the candidate is actually low quality, it obfuscates by generating signal σ^l with a smaller probability. The overall effect on the informativeness of signal is generally ambiguous, and can depend on the prior likelihood ratio p_H/p_L.

ChatGPT generally also has ambiguous effects on the perceived qualities of candidates with signal σ^h and signal σ^l. We have

$$\mathrm{d}\beta_H^h = \beta_H^h(1 - \beta_H^h)\,\mathrm{d}\ln\left(\frac{q_H^h}{q_L^h}\right),$$

and

$$\mathrm{d}\beta_H^l = \beta_H^l(1-\beta_H^l)\,\mathrm{d}\ln\left(\frac{1-q_H^h}{1-q_L^h}\right),$$

with

$$\mathrm{d}u^j = (v_H - v_L)\mathrm{d}\beta_H^j$$

for each $j = h, l$. Thus, ChatGPT can either increase or decrease the perceived quality u^h of a candidate with signal σ^h, depending whether it increases or decreases the likelihood ratio q_H^h/q_L^h of actual high-quality candidates with signal σ^h compared to actual low-quality candidates. Similarly, ChatGPT can either increase or decrease the perceived quality u^h of a candidate with signal σ^l, depending whether it increases or decreases the likelihood ratio $(1-q_H^h)/(1-q_L^h)$ of actual high-quality candidates with signal σ^l compared to actual low-quality candidates. However, since $q_H^h > q_L^h$, we have

$$\frac{q_H^h}{q_L^h} > \frac{1-q_H^h}{1-q_L^h}.$$

Thus, under ChatGPT, if $\mathrm{d}u^h \geq 0$ then $\mathrm{d}u^l < 0$, and if $\mathrm{d}u^l \geq 0$ then $\mathrm{d}u^h < 0$. The use of ChatGPT cannot simultaneously increase the posterior means of both signals σ^h and σ^l. However, the two posterior means can both decrease. In particular, if

$$\frac{1-q_H^h}{1-q_L^h} < \frac{\mathrm{d}q_H^h}{\mathrm{d}q_L^h} < \frac{q_H^h}{q_L^h}.$$

then $\mathrm{d}u^j < 0$ for each $j = h, l$.

We can categorize ChatGPT into three cases depending on the magnitude of $\mathrm{d}q_H^h/\mathrm{d}q_L^h$. In the case of $\mathrm{d}u^h \geq 0 > \mathrm{d}u^l$, the spread between the two posterior means, u^h and u^l, is increased. This is the case where the simultaneous increases in q_H^h and in $q_L h$ increase the likelihood ratio q_H^h/q_L^h. In the opposite case of $\mathrm{d}u^h < 0 \leq \mathrm{d}u^l$, the spread is decreased as the simultaneous increases in q_H^h and in $q_L h$ increase the likelihood ratio $(1-q_H^h)/(1-q_L^h)$. The following lemma uses the entropy of V conditional on the signal σ, which measures the expected residual uncertainty about V after σ is realized, to verify that, regardless of the prior likelihood ratio p_H/p_L, signals about the underlying quality of candidates become more "informative on average" in the first case but less informative on average in the second case. This result is intuitive, because the extreme case of $\mathrm{d}q_H^h > 0$ and $\mathrm{d}q_L^h = 0$ should make signals more informative and the opposite case of $\mathrm{d}q_H^h = 0$ and $\mathrm{d}q_L^h > 0$ should make signals less informative. The conditional entropy is given by

$$S(V|\sigma) = -p_H(q_H^h \ln \beta_H^h + (1-q_H^h)\ln \beta_H^l) - p_L(q_L^h \ln(1-\beta_H^h) + (1-q_L^h)\ln(1-\beta_H^l)).$$

Lemma 1. *For any prior likelihood ratio p_H/p_L, under ChatGPT, $du^h \geq 0 > du^l$ and the conditional entropy decreases if $dq_H^h/dq_L^h \geq q_H^h/q_L^h$, and $du^h < 0 \leq du^l$ and the conditional entropy increases if $dq_H^h/dq_L^h \leq (1-q_H^h)/(1-q_L^h)$.*

The proof of this lemma is in the appendix. We refer to the case of $\mathrm{d}u^h \geq 0 > \mathrm{d}u^l$ as the "informative ChatGPT," and the case of $\mathrm{d}u^h < 0 \leq \mathrm{d}u^l$ as the "uninformative ChatGPT." These two cases highlight how the use of ChatGPT can unambiguously increase or decrease the average informativeness of job market papers as a signal about quality of candidates, even though in both cases it makes signal σ^h more likely regardless of the underlying quality of a candidate.

In the "middle" case of $\mathrm{d}u^h < 0$ and $\mathrm{d}u^l < 0$, whether the average informativeness increases or decreases depends on the prior likelihood ratio p_H/p_L. We refer this case as "ambiguous ChatGPT." From the proof of Lemma 1, the conditional entropy decreases if and only if

$$\frac{\mathrm{d}q_H^h}{\mathrm{d}q_L^h} > \frac{p_L}{p_H} \cdot \frac{\ln(\pi^h/(1-\pi^h)) - \ln(q_L^h/(1-q_L^h))}{\ln(q_H^h/(1-q_H^h)) - \ln(\pi^h/(1-\pi^h))} \equiv \mu_p,$$

with

$$\frac{1-q_H^h}{1-q_L^h} < \mu_p < \frac{q_H^h}{q_L^h}.$$

for all p_H/p_L. The critical value μ_p of $\mathrm{d}q_H^h/\mathrm{d}q_L^h$ that keeps the conditional entropy unchanged under ambiguous ChatGPT depends on the prior likelihood ratio p_H/p_L.

Lemma 1 provides a characterization of the link between the effect of ChatGPT on the average informativeness of signal and its effect on the ratio of the posterior means u^h/u^l. Under informative ChatGPT, the ratio u^h/u^l increases because $\mathrm{d}u^h \geq 0 > \mathrm{d}u^l$; under uninformative ChatGPT, the ratio u^h/u^l decreases because $\mathrm{d}u^h < 0 \leq \mathrm{d}u^l$. Under ambiguous ChatGPT, the ratio increases if and only if

$$\frac{\beta_H^h(1-\beta_H^h)}{u^h}\mathrm{d}\ln\left(\frac{q_H^h}{q_L^h}\right) > \frac{\beta_H^l(1-\beta_H^l)}{u^l}\mathrm{d}\ln\left(\frac{1-q_H^h}{1-q_L^h}\right).$$

The terms on the two sides of the above inequality are both negative, because $du^j < 0$ for each $j = h, l$ under the the ambiguous ChatGPT. The above is equivalent to

$$\frac{\mathrm{d}q_H^h}{\mathrm{d}q_L^h} > \frac{u^h(1-q_H^h)(\pi^h)^2 + u^l q_H^h (\pi^l)^2}{u^h(1-q_L^h)(\pi^h)^2 + u^l q_L^h (\pi^l)^2} \equiv \mu_v,$$

with

$$\frac{1-q_H^h}{1-q_L^h} < \mu_v < \frac{q_H^h}{q_L^h}.$$

Unlike μ_p, the critical value of $\mathrm{d}q_H^h/\mathrm{d}q_L^h$ that keeps the ratio u^h/u^l unchanged under ambiguous ChatGPT depends also on the ratio v_H/v_L. It can be verified that μ_v is decreasing in v_H/v_L.

3 Symmetric Equilibrium

Fix any number of candidates $\gamma^h \geq 1$ with signal σ^h and $\gamma^l \geq 1$ with signal σ^l, with $\gamma^h + \gamma^l = C$. The probability of this realization is given by a binomial distribution

$$\frac{C!}{\gamma^h!\gamma^l!}(\pi^h)^{\gamma^h}(\pi^l)^{\gamma^l}.$$

Since V_c is independent across C candidates, the numbers γ^h and γ^l are uninformative about the underlying quality V_c of any individual candidate.

Following Peters (2010), we can show that there is a unique symmetric mixed-strategy equilibrium of the game of the following form. There is a threshold offer $\hat{x} \in [0, \overline{x}]$, and a mixing probability $w \in (0, 1)$, such that the probability that any department d, $d = 1, \ldots, D$, makes an offer to a randomly selected candidate with signal σ^h is 1 if $x_d \geq \hat{x}$ and w if $x_d < \hat{x}$.

The threshold type $\hat{x}$ is determined by the indifference condition:

$$u^h \left(1 - \frac{1}{\gamma^h}(1 - F(\hat{x}))\right)^{D-1} = u^l.$$

The left-hand side is u^h times the probability that an offer by type $\hat{x}$ to a randomly selected candidate with signal σ^h is accepted, given by

$$\sum_{\delta(\hat{x},\overline{x})=0}^{D-1} \frac{(D-1)!}{\delta(\hat{x},\overline{x})!(D-1-\delta(\hat{x},\overline{x}))!}(1 - F(\hat{x}))^{\delta(\hat{x},\overline{x})} \left(1 - \frac{1}{\gamma^h}\right)^{\delta(\hat{x},\overline{x})} (F(\hat{x}))^{D-1-\delta(\hat{x},\overline{x})},$$

where, for any interval $[x_1, x_2] \subseteq [0, \overline{x}]$, $\delta(x_1, x_2)$ is the realized number of departments with offers from the interval. That is, type $\hat{x}$ is indifferent between making an offer to a randomly selected candidate with signal σ^h, which is accepted when all departments with offers higher than $\hat{x}$ make their offers to one of other $\gamma^h - 1$ candidates with signal σ^h, and making an offer to a randomly selected candidate with signal σ^l, which is always accepted. To save notation, define

$$\rho = \left(\frac{u^h}{u^l}\right)^{1/(D-1)}.$$

We assume that

$$\frac{C}{C-1} > \rho,$$

so that there is a unique interior $\hat{x} \in (0, \overline{x})$ that satisfies the indifference condition regardless of the realized number γ^h of graduations with signal σ^h (so long as $1 \leq \gamma^h \leq C$).

For any department d with offer $x_d > \hat{x}$, it is strictly preferred to make an offer to a randomly selected candidate with signal σ^h.

Fix any department d with offer $x_d < \hat{x}$, the following indifference condition holds:

$$u^h \left(1 - \frac{1}{\gamma^h}(1 - F(\hat{x})) - \frac{1}{\gamma^h}(F(\hat{x}) - F(x_d))w\right)^{D-1} = u^l \left(1 - \frac{1}{\gamma^l}(F(\hat{x}) - F(x_d))(1-w)\right)^{D-1}.$$

On the left-hand side of the above condition, the probability that a randomly selected candidate with signal σ^h accepts the offer x_d is equal to the probability that all offers from rival departments with higher offers than x_d go to one of other candidates with signal σ^h. This can be computed as:

$$\sum_{\delta(x_d,\hat{x})+\delta(\hat{x},\overline{x})=0}^{D-1} \frac{(D-1)!}{\delta(\hat{x},\overline{x})!\delta(x_d,\hat{x})!(D-1-\delta(\hat{x},\overline{x})-\delta(x_d,\hat{x}))!}(F(x_d))^{D-1-\delta(\hat{x},\overline{x})-\delta(x_d,\hat{x})}$$
$$\cdot (1-F(\hat{x}))^{\delta(\hat{x},\overline{x})}\left(1-\frac{1}{\gamma^h}\right)^{\delta(\hat{x},\overline{x})}(F(\hat{x})-F(x_d))^{\delta(x_d,\hat{x})}\left(1-\frac{w}{\gamma^h}\right)^{\delta(x_d,\hat{x})}.$$

On the right-hand side of the above condition, the probability that a randomly selected candidate with signal σ^l accepts the offer x_d is given by

$$\sum_{\delta(x_d,\hat{x})+\delta(\hat{x},\overline{x})=0}^{D-1} \frac{(D-1)!}{\delta(\hat{x},\overline{x})!\delta(x_d,\hat{x})!(D-1-\delta(\hat{x},\overline{x})-\delta(x_d,\hat{x}))!}(1-F(\hat{x}))^{\delta(\hat{x},\overline{x})}(F(x_d))^{D-1-\delta(\hat{x},\overline{x})-\delta(x_d,\hat{x})}$$
$$\cdot (F(\hat{x})-F(x_d))^{\delta(x_d,\hat{x})}\left(1-\frac{1}{\gamma^l}(1-w)\right)^{\delta(x_d,\hat{x})}.$$

Combining the indifference condition of type $x_d < \hat{x}$ and the indifference condition for type $x_d = \hat{x}$, we have

$$w = \frac{\gamma^h}{\gamma^h + \gamma^l \rho}.$$

Thus, there is a unique interior value $w \in (0,1)$ for any realization $\gamma^h \geq 1$ and $\gamma^l \geq 1$. Since w is independent of the offer x_d, we have verified the existence of a unique symmetric equilibrium at the interim stage.

Lemma 2. *Fix any number of candidates $\gamma^h \geq 1$ with signal σ^h and $\gamma^l \geq 1$ with signal σ^l, with $\gamma^h + \gamma^l = C$. There exists a unique symmetric equilibrium, with $\hat{x} \in (0, \overline{x})$ and $w \in (0, 1)$, such that the probability of type x makes an offer to a candidate with signal σ^h is 1 if $x \geq \hat{x}$ and w if $x < \hat{x}$.*

3.1 Comparative statics

Fix any number of candidates $\gamma^h \geq 1$ with signal σ^h and $\gamma^l \geq 1$ with signal σ^l, with $\gamma^h + \gamma^l = C$. At the interim stage, the symmetric equilibrium with threshold type $\hat{x}$ and probability w of types below $\hat{x}$ targeting a candidate with signal σ^h are completely determined by the posterior means u^h and u^l.

From the indifference condition of the threshold type, the value of $\hat{x}$ depends on the ratio of u^l to u^h. An increase in ρ reduces $\hat{x}$. This is intuitive. When candidates with signal σ^h become more attractive to departments relative to those with signal σ^l, the threshold type $\hat{x}$ before the change strictly prefers targeting candidates with signal σ^h to targeting candidates with signal σ^l. The threshold has to decrease in order to restore the indifference.

The equilibrium value of w also depends only on the ratio of u^l to u^h. From the expression for w, an increase in ρ reduces the equilibrium probability w that departments with offers below the threshold target candidates with signal σ^h. This is perhaps unintuitive. After all, candidates with signal σ^h have become more attractive to departments relative to those with signal σ^l. The reason that w decreases is that any type $x < \hat{x}$ faces the same posterior means u^h and u^l as the threshold type $\hat{x}$, but while type $\hat{x}$ loses only to types above $\hat{x}$ in making an offer to a candidate with signal σ^h, type x also loses to types between x and $\hat{x}$ that target these candidates with probability w. When ρ increases, w has to decrease to keep type x also indifferent.

Fix any $x \leq \hat{x}$. For welfare effects of changes in signal accuracies on departments and candidates, we also need comparative statics on

$$A^h(x) = \frac{1}{\gamma^h}(1 - F(\hat{x})) + \frac{1}{\gamma^h}(F(\hat{x}) - F(x))w,$$

which is the probability of a given candidate with signal σ^h receiving an offer of x or above from a randomly selected department, and on

$$A^l(x) = \frac{1}{\gamma^l}(F(\hat{x}) - F(x))(1 - w),$$

which is the probability of a given candidate with signal σ^l receiving an offer of x or above from a randomly selected department. By indifference condition of type x, we have

$$\frac{1 - A^l(x)}{1 - A^h(x)} = \rho$$

for all $x \leq \hat{x}$. At the same time

$$\gamma^h A^h(x) + \gamma^l A^l(x) = 1 - F(x),$$

which is independent of ρ. Thus, as ρ increases, $A^h(x)$ increases and $A^l(x)$ decreases for all $x \leq \hat{x}$.

For any $x > \hat{x}$, the probability of a given candidate with signal σ^h receiving an offer of x or above from a randomly selected department is

$$A^h(x) = \frac{1}{\gamma^h}(1 - F(x)),$$

and the probability of a given candidate with signal σ^h receiving an offer of x or above from a randomly selected department is

$$A^l(x) = 0.$$

Both $A^h(x)$ and $A^l(x)$ are unaffected by changes in ρ.

We summarize the comparative statics as follows.

Proposition 1. *As ρ increases, for any realized $\gamma^h \geq 1$ and $\gamma^l \geq 1$, with $\gamma^h + \gamma^l = C$, both $\hat{x}$ and w decrease; $A^h(x)$ and $A^l(x)$ remain unchanged for $x > \hat{x}$, and $A^h(x)$ increases and $A^l(x)$ decreases for $x \leq \hat{x}$.*

The above result implies that the two "unambiguous" cases of ChatGPT have opposite comparative statics. Under the informative ChatGPT, more high offers and fewer low offers go to candidates with signal σ^h; under the uninformative ChatGPT, fewer high offers and more low offers go to candidates with signal σ^h. In ambiguous case, the comparative statics of the use of ChatGPT on the symmetric equilibrium depends only on its effect on the ratio ρ, or equivalently, u^h/u^l. Thus, ρ increases if and only if $\mathrm{d}q_H^h/\mathrm{d}q_L^h > \mu_v$.

3.2 Candidates welfare

Fix any number of candidates $\gamma^h \geq 1$ with signal σ^h and $\gamma^l \geq 1$ with signal σ^l, with $\gamma^h + \gamma^l = C$. For any candidate with signal σ^l, the match rate $M^l(x)$ is defined as the probability of receiving at least one offer from a department d with offer $x_d \geq x$. By equilibrium construction, $M^l(x) = 0$ for any $x \geq \hat{x}$, and for any $x < \hat{x}$, $M^l(x)$ is given by

$$\sum_{\delta(x,\hat{x})+\delta(\hat{x},\overline{x})=0}^{D} \frac{D!}{\delta(\hat{x},\overline{x})!\delta(x,\hat{x})!(D-\delta(\hat{x},\overline{x})-\delta(x,\hat{x}))!}(1-F(\hat{x}))^{\delta(\hat{x},\overline{x})}(F(x))^{D-\delta(\hat{x},\overline{x})-\delta(x,\hat{x})}$$
$$\cdot (F(\hat{x})-F(x))^{\delta(x,\hat{x})} \sum_{\delta=0}^{\delta(x,\hat{x})} \frac{\delta(x,\hat{x})!}{\delta!(\delta(x,\hat{x})-\delta)!} w^{\delta(x,\hat{x})-\delta}(1-w)^{\delta}\left(1-\left(1-\frac{1}{\gamma^l}\right)^{\delta}\right)$$

Thus,

$$M^l(x) = 1-\left(1-\frac{1}{\gamma^l}(F(\hat{x})-F(x))(1-w)\right)^D = 1-(1-A^l(x))^D.$$

Now, consider the match rate $M^h(x)$ of a candidate with signal σ^h, defined as the probability of receiving at least one offer from a department d with offer $x_d \geq x$. For any

$x \geq \hat{x}$, $M^h(x)$ is given by

$$M^h(x) = \sum_{\delta(x,\overline{x})=0}^{D} \frac{D!}{\delta(x,\overline{x})!(D-\delta(x,\overline{x}))!} (F(x))^{D-\delta(x,\overline{x})} (1-F(x))^{\delta(x,\overline{x})} \left(1 - \left(1 - \frac{1}{\gamma^h}\right)^{\delta(x,\overline{x})}\right).$$

Thus,

$$M^h(x) = 1 - \left(1 - \frac{1}{\gamma^h}(1-F(x))\right)^D = 1 - (1 - A^h(x))^D.$$

For $x < \hat{x}$, offers to a candidate with signal σ^h can also come from randomization by departments with offers between x and $\hat{x}$. The match rate $M^h(x)$ can be obtained by calculating the probability no offer is made by these randomizing departments or by any department with offers higher than $\hat{x}$. This probability is give by

$$\sum_{\delta(x,\hat{x})+\delta(\hat{x},\overline{x})=0}^{D} \frac{D!}{\delta(\hat{x},\overline{x})!\delta(x,\hat{x})!(D-\delta(\hat{x},\overline{x})-\delta(x,\hat{x}))!} (1-F(\hat{x}))^{\delta(\hat{x},\overline{x})} \left(1 - \frac{1}{\gamma^h}\right)^{\delta(x,\overline{x})}$$
$$\cdot (F(\hat{x}) - F(x))^{\delta(x,\hat{x})} \sum_{\delta=0}^{\delta(x,\hat{x})} \frac{\delta(x,\hat{x})!}{\delta!(\delta(x,\hat{x})-\delta)!} (1-w)^{\delta(x,\hat{x})-\delta} w^{\delta} \left(1 - \frac{1}{\gamma^h}\right)^{\delta} (F(x))^{D-\delta(\hat{x},\overline{x})-\delta(x,\hat{x})}$$

Thus,

$$M^h(x) = 1 - \left(1 - \frac{1}{\gamma^h}(1-F(\hat{x})) - \frac{1}{\gamma^h}(F(\hat{x}) - F(x))w\right)^D = 1 - (1 - A^h(x))^D.$$

The welfare of candidates depends not only on the total match rates $M^l(0)$ and $M^h(0)$, but also on the entire match rate functions $M^l(\cdot)$ and $M^h(\cdot)$. More precisely, the payoff to a candidate with signal σ^j, $j = h, l$ is given by

$$R^j = -\int_0^{\overline{x}} x \frac{\mathrm{d}M^j(x)}{\mathrm{d}x} \mathrm{d}x.$$

Since $M^j(\overline{x}) = 0$, it follows from integration by parts that

$$R^j = \int_0^{\overline{x}} M^j(x) \mathrm{d}x.$$

By Proposition 1, as ρ increases, the match rate function $M^h(x)$ for candidates with signal σ^h stays the same for $x \geq \hat{x}$ and increases for all $x < \hat{x}$. The opposite is true for candidates with signal σ^l: $M^l(x)$ stays 0 for all $x \geq \hat{x}$ but decreases for all $x > \hat{x}$. Thus, we have the following comparative statics on interim candidate welfare.

Proposition 2. *For any realized $\gamma^h \geq 1$ and $\gamma^l \geq 1$, with $\gamma^h + \gamma^l = C$, candidates with high signals are better off and those with low signals are worse off if ρ increases, and candidates with high signals are worse off and those with low signals are better off if ρ decreases.*

By Proposition 2, ChatGPT has exactly opposite effects on the welfare of candidates with signal σ^h and signal σ^l. If ChatGPT is informative, candidates with signal σ^h are better off and those with signal σ^l are worse off. The opposite is true if ChatGPT is uninformative: candidates with signal σ^h are worse off and those with signal σ^l are better off. When ChatGPT is ambiguous, the effect depends only on whether ChatGPT increases or decreases ρ, that is, whether or not $\mathrm{d}q_H^h/\mathrm{d}q_L^h$ is greater than μ_v.

To investigate the impact of ChatGPT on the ex ante welfare of candidates, define R_i, $i = H, L$, as the expected equilibrium payoff of a candidate whose underlying quality is $V = v_i$. Then

$$R_i = q_i^h R^h + (1 - q_i^h) R^l.$$

We can then define the ex ante equilibrium payoff of candidates as

$$R = p_H R_H + p_L R_L = \pi^h R^h + \pi^l R^l.$$

Since $A^h(x) > A^l(x)$ for all $x \in (0, \overline{x})$, and thus $R^h > R^l$, and $\mathrm{d}\pi^h > 0$ under ChatGPT, the ex ante impact is positive. The equilibrium impact in general can be positive or negative because, by Proposition 2, the interim payoffs R^h and R^l are changed in opposite ways. However, the equilibrium effect is nil when

$$\frac{\mathrm{d}q_H^h}{\mathrm{d}q_L^h} = \mu_v.$$

By continuity, we have the following result.

Corollary 1. *There exists an interval of dq_H^h/dq_L^h such that ambiguous ChatGPT makes all candidates ex ante better off.*

When ChatGPT is ambiguous, its effect on the average informativeness of signals depends on the prior likelihood ratio. This means the equilibrium effect on ρ is relatively small. Since ChatGPT generates signal σ^h more often regardless of the underlying quality, all candidates are better off ex ante.

3.3 Departments welfare

Fix any number of candidates $\gamma^h \geq 1$ with signal σ^h and $\gamma^l \geq 1$ with signal σ^l, with $\gamma^h + \gamma^l = C$. The equilibrium payoff for a department d with offer $x_d \geq \hat{x}$ is given by

$$P(x_d) = u^h \left(1 - \frac{1}{\gamma^h}(1 - F(x_d))\right)^{D-1}.$$

To get an expression for a department d with offer $x_d < \hat{x}$, we solve the indifference condition to get

$$P(x_d) = \left(\frac{\gamma^h + \gamma^l - (1 - F(x_d))}{\gamma^h (u^h)^{-1/(D-1)} + \gamma^l (u^l)^{-1/(D-1)}}\right)^{D-1} = u^h \left(\frac{\gamma^h + \gamma^l - (1 - F(x_d))}{\gamma^h + \gamma^l \rho}\right)^{D-1}.$$

The equilibrium payoff $P(x_d)$ depends on the offer x_d only through $F(x_d)$. There is a kink at $x_d = \hat{x}$, the highest offer such that the department remains indifferent between making an offer to a candidate with signal σ^h and making an offer to a candidate with signal σ^l. The kink is "convex," reflecting the fact that departments with offers slightly higher than the threshold $\hat{x}$ exclusively target γ^h candidates with signal σ^h, while those with offers slightly below $\hat{x}$ mix between γ^h candidates with signal σ^h and γ^l candidates with signal σ^l.

Since departments with offers above the threshold strictly prefer to make an offer to candidates with signal σ^h, they are better off if and only if u^h becomes higher. The sign of the impact on all departments with offers below the threshold is independent of their offer. We have the following result immediately.

Proposition 3. *For any realized $\gamma^h \geq 1$ and $\gamma^l \geq 1$, with $\gamma^h + \gamma^l = C$, departments with offers above $\hat{x}$ are better off if and only if u^h increases, and those with offers below $\hat{x}$ are better off if and only if $\gamma^h (u^h)^{-1/(D-1)} + \gamma^l (u^l)^{-1/(D-1)}$ decreases.*

By Proposition 3, the welfare impacts of ChatGPT on departments generally depend on whether ChatGPT is informative, ambiguous, or uninformative, and depend on whether departments have offers above the threshold or below the threshold Departments with offers above the threshold are better off if ChatGPT is informative, and are otherwise worse off. For departments with offers below the threshold, they are worse off if ChatGPT is ambiguous, and by continuity, are worse off so long as ChatGPT is not too informative or too uninformative, that is, if the ratio $\mathrm{d}q^h_H/\mathrm{d}q^h_L$ is not extreme.

Indeed, the impact of ChatGPT on the welfare of departments with offers below the threshold has a monotone relationship with the magnitude of $\mathrm{d}q^h_H/\mathrm{d}q^h_L$. By Proposition 3, the sign of $\mathrm{d}P(x_d)$ for any department d with offer $x_d < \hat{x}$ is the same as

$$\gamma^h (u^h)^{-D/(D-1)} \mathrm{d}u^h + \gamma^l (u^l)^{-D/(D-1)} \mathrm{d}u^l,$$

where

$$\mathrm{d}u^h = \frac{p_H p_L (v_H - v_L)}{(\pi^h)^2} (q^h_L \mathrm{d}q^h_H - q^h_H \mathrm{d}q^h_L),$$

and

$$\mathrm{d}u^l = \frac{p_H p_L (v_H - v_L)}{(\pi^l)^2} ((1 - q^h_H) \mathrm{d}q^h_L - (1 - q^h_L) \mathrm{d}q^h_H).$$

For informative ChatGPT, we have

$$\frac{\mathrm{d}q^h_H}{\mathrm{d}q^h_L} \geq \frac{q^h_H}{q^h_L} > \frac{1 - q^h_H}{1 - q^h_L},$$

and so an increase in $\mathrm{d}q^h_H/\mathrm{d}q^h_L$ makes it more likely that $\mathrm{d}P(x_d) > 0$. That is, a very informative ChatGPT may make departments with offers below the threshold better off by increasing u^h, despite of decreasing u^l at the same time. Conversely, for uninformative ChatGPT, we have

$$\frac{\mathrm{d}q_H^h}{\mathrm{d}q_L^h} \leq \frac{1-q_H^h}{1-q_L^h} < \frac{q_H^h}{q_L^h},$$

and so a decrease in $\mathrm{d}q_H^h/\mathrm{d}q_L^h$ makes it more likely that $\mathrm{d}P(x_d) > 0$. A very uninformative ChatGPT may make departments with offers below the threshold better off by increasing u^l, despite of decreasing u^h at the same time.

From the ex ante perspective, the payoff of a department d with offer x^d depends on the realization of number of candidates γ^h and γ^l with signal σ^h and signal σ^l respectively. This payoff is given by

$$\sum_{\gamma^h=0}^{C} \frac{C!}{\gamma^h!(C-\gamma^h)!} (\pi^h)^{\gamma^h} (1-\pi^h)^{C-\gamma^h} P(x_d),$$

where

$$P(x_d) = u^l \left(1 - \frac{1}{C}(1-F(x_d))\right)^{D-1}$$

for $\gamma^h = 0$, and

$$P(x_d) = u^h \left(1 - \frac{1}{C}(1-F(x_d))\right)^{D-1}$$

for $\gamma^h = C$. By increasing π^h, ChaptGPT shifts the binomial distribution of γ^h and γ^l towards signal σ^h in the sense that the likelihood ratio

$$\frac{(\pi^h + \mathrm{d}\pi^h)^{\gamma^h}(1-\pi^h-\mathrm{d}\pi^h)^{C-\gamma^h}}{(\pi^h)^{\gamma^h}(1-\pi^h)^{C-\gamma^h}}$$

increases with γ^h. By Proposition 3, the ex ante effect on departments is more likely to be positive if ChatGPT increases u^h, and negative if ChatGPT decreases u^h. Combining with the earlier analysis of the equilibrium effects of ChatGPT, we see that very informative ChatGPT is more likely to make departments better off ex ante than very uninformative

ChatGPT. When ChatGPT is ambiguous, the equilibrium effects are always negative. In this case, the sign of the overall effect no longer depends on the realization of γ^h and γ^l. We have the following corollary immediately.

Corollary 2. *Ambiguous ChatGPT makes all departments ex ante worse off.*

When ChatGPT makes the average informativeness of signals unambiguously higher or lower, the equilibrium effects of ChatGPT is generally ambiguous because it causes the perceived mean quality u^h and the perceived mean quality u^l go in opposite directions. When its effect on the average informativeness of signals depends on the prior likelihood ratio, ChatGPT reduces the two perceived means simultaneously. At the interim stage, all departments are worse off, whether they exclusively target candidates with signal σ^h or randomize between these candidates and those with signal σ^l. Although the ex ante effect of ChatGPT favors departments that target candidates with signals σ^h, even these departments are worse off ex ante.

4 Discussion

In classical information theory, Blackwell (1953) develops a general notion of comparing information structures for all decision problems, and Lehmann's (1988) ordering of information structures applies to monotone decision problems. The premise of this paper is that the use of ChatGPT makes job market papers appear well-written regardless of their true qualities. The impact of ChatGPT on the informativeness of job market papers falls outside the existing analytical frameworks. To make progress, we have modeled ChatGPT in the simplest information environment of two states and two signals, and use conditional entropy to quantify the impact. The question is whether we can extend the model to more than two states and more than two signals.

Consider the following extension. There are I states that are decreasingly ordered $v_1 > \ldots > v_I$, with prior beliefs $p_H, \ldots, p_I$, and J signals, $\sigma^1, \ldots, \sigma^J$. Let

$$q_i^j = \Pr(\sigma^j | v_i)$$

for each $i = 1, \ldots, I$ and $j = 1, \ldots, J$. Assume that monotone likelihood ratio property holds – the ratio $q_i^j / q_i^{j'}$ is decreasing in i for all $j < j'$ – so that a lower-indexed signal is more likely associated with a lower-indexed state. For each $j = 1, \ldots, J$, let

$$\pi^j = \sum_{i=1}^{I} p_i q_i^j$$

be the total probability that σ_j is realized, let

$$\beta_i^j = \Pr(v_i | \sigma^j) = \frac{p_i q_i^j}{\pi^j}$$

be the posterior probability that the state is v_i conditional on signal σ^j, and let

$$u^j = \sum_{i=1}^{I} \beta_i^j v_i$$

be the posterior mean state conditional on signal σ^j. By the monotone likelihood ratio property, $\beta_i^j / \beta_i^{j'}$ is decreasing in i for all $j < j'$, implying u^j is decreasing in j – posterior means are higher for lower-indexed signals. To model the use of ChatGPT, let $\mathrm{d}q_i^j$ be the infinitesimal change in q_i^j, and suppose that for each $i = 1, \ldots, I$, and each $j = 1, \ldots, J$,

$$\sum_{j'=1}^{j} \mathrm{d}q_i^{j'} \geq 0.$$

so that there is a first-order stochastic dominant shift towards lower-indexed signals.

Consider the special case of the above framework with only two signals, $j = h, l$. By the monotone likelihood ratio property, q_i^h is greater for lower-indexed states. The conditional entropy is given by

$$S(V|\sigma) = \sum_{i=1}^{I} p_i (q_i^h \ln \beta_i^h + (1 - q_i^h) \ln \beta_i^h).$$

Using

$$\sum_{i=1}^{I} \mathrm{d}\beta_i^h = \sum_{i=1}^{I} \mathrm{d}\beta_i^l = 0,$$

we have

$$\mathrm{d}S(V|\sigma) = \sum_{i=1}^{I} p_i \ln\left(\frac{\beta_i^h}{\beta_i^l}\right) \mathrm{d}q_i^h.$$

The above takes the same form as in the case of two states. For changes in posterior means caused by the use of ChatGPT, for each $j = h, l$, we have

$$\mathrm{d}u^j = \sum_{i=1}^{I} \mathrm{d}\beta_i^j v_i,$$

where

$$\frac{\mathrm{d}\beta_i^j}{\beta_i^j} = \frac{\mathrm{d}q_i^j}{q_i^j} - \frac{\mathrm{d}\pi^j}{\pi^j}.$$

Unlike in the case of two states, it is not possible to state the conditions for the signs of $\mathrm{d}u^h$ and $\mathrm{d}u^l$ independently of the states and prior beliefs about the states. Nonetheless, the law of iterated expectations requires

$$\pi^h u^h + \pi^l u^l = \sum_{i=1}^{I} p_i v_i,$$

and so

$$(u^h - u^l)\mathrm{d}\pi^h + \pi^h \mathrm{d}u^h + \pi^l \mathrm{d}u^l = 0.$$

Since $u^h > u^l$ and $\mathrm{d}\pi^h > 0$, as in the case of two states, we have either $\mathrm{d}u^h \geq 0 > \mathrm{d}u^l$, or $\mathrm{d}u^h < 0 \leq \mathrm{d}u^l$, or $\mathrm{d}u^h < 0$ and $\mathrm{d}u^l < 0$. We conjecture, as an extension of the characterization result of Lemma 1, that ChatGPT improves the conditional entropy for prior beliefs if it increases the perceived quality u^h upon the signal σ^h, and reduces the conditional entropy if it increases u^l upon σ^l.

Appendix: Proof of Lemma 1

Denote as $S(V|\sigma)$ the entropy of the random variable V conditional on σ. We have

$$\begin{aligned}
\mathrm{d}S(V|\sigma) = & -p_H\left(\ln\left(\frac{\beta_H^h}{\beta_H^l}\right)\mathrm{d}q_H^h + \frac{q_H^h}{\beta_H^h}\mathrm{d}\beta_H^h + \frac{1-q_H^h}{\beta_H^l}\mathrm{d}\beta_H^l\right) \\
& - p_L\left(\ln\left(\frac{1-\beta_H^h}{1-\beta_H^l}\right)\mathrm{d}q_L^h - \frac{q_L^h}{1-\beta_H^h}\mathrm{d}\beta_H^h - \frac{1-q_L^h}{1-\beta_H^l}\mathrm{d}\beta_H^l\right) \\
= & -p_H\ln\left(\frac{\beta_H^h}{\beta_H^l}\right)\mathrm{d}q_H^h - p_L\ln\left(\frac{1-\beta_H^h}{1-\beta_H^l}\right)\mathrm{d}q_L^h.
\end{aligned}$$

First, suppose that

$$\frac{\mathrm{d}q_H^h}{\mathrm{d}q_L^h} \geq \frac{q_H^h}{q_L^h},$$

so that $\mathrm{d}u^h \geq 0 > \mathrm{d}u^l$. We claim that the sign of $\mathrm{d}S(V|\sigma)$ is negative for all p_H/p_L. Since $\beta_H^h > \beta_H^l$, this is equivalent to showing that

$$-\frac{\beta_H^h}{1-\beta_H^h}\ln\left(\frac{\beta_H^h}{\beta_H^l}\right) - \ln\left(\frac{1-\beta_H^h}{1-\beta_H^l}\right) < 0.$$

Define

$$K \equiv \frac{\beta_H^h}{1-\beta_H^h}\frac{1-\beta_H^l}{\beta_H^l} = \frac{q_H^h}{q_L^h}\frac{1-q_L^h}{1-q_H^h},$$

which is independent of p_H/p_L, and is strictly greater than 1 because $\beta_H^h > \beta_H^l$. We can rewrite the desired inequality as

$$-\frac{\ln(1+(K-1)(1-\beta_H^h))}{1-\beta_H^h} + \ln K < 0.$$

For fixed q_H^h and q_L^h, we can take the left-hand side of the above inequality as a function of β_H^h when p_H/p_L varies for fixed K. By taking derivatives we can easily show that the function $\ln(1+z)/z$ is strictly decreasing for all $z > 0$. Since $K > 1$, the desired inequality holds for all $\beta^h_H > 0$, and hence all p_H/p_L. We have established the claim that $\mathrm{d}S(V|\sigma) < 0$ for all p_H/p_L.

Second, suppose that

$$\frac{\mathrm{d}q^h_H}{\mathrm{d}q^h_L} \leq \frac{1-q^h_H}{1-q^h_L},$$

so that $\mathrm{d}u^h < 0 \leq \mathrm{d}u^l$. We claim that $\mathrm{d}S(V|\sigma) > 0$ for all p_H/p_L. Since $\beta^h_H > \beta^l_H$, this is equivalent to showing that

$$-\ln\left(\frac{\beta^h_H}{\beta^l_H}\right) - \frac{1-\beta^l_H}{\beta^l_H}\ln\left(\frac{1-\beta^h_H}{1-\beta^l_H}\right) > 0,$$

or equivalently,

$$-\ln K + \frac{\ln(1+(K-1)\beta^l_H)}{\beta^l_H} > 0.$$

Since $\ln(1+z)/z$ is strictly decreasing for all $z > 0$, the above inequality holds for all $\beta^l_H < 1$, and hence for all p_H/p_L. We have established the claim that $\mathrm{d}S(V|\sigma) > 0$ for all p_H/p_L.

References

[1] David Blackwell, Equivalent Comparisons of Experiments, The Annals of Mathematical Statistics, 24(2), pp. 265-272, 1953.

[2] Erich Lehmann, Comparing Location Experiments, Annals of Statistics, 16(2), pp. 521-533, 1988.

[3] Michael Peters, Noncontractible Heterogeneity in Directed Search, Econometrica, 78(4), pp. 1173-1200, 2010.