Employee Recruiting and the Lake Wobegon Effect*

John Moran
Department of Economics
Syracuse University

John Morgan
Woodrow Wilson School
Princeton University

June 2001

Abstract

Employers, educational institutions, and other organizations are often faced with the problem of selecting the most qualified candidate to fill an available position. To this end, many employers have adopted a tournament-like procedure consisting of an initial phase in which recommendations solicited from third-party "referees" are used to eliminate unqualified candidates, followed by an interview phase in which the remaining candidates are ordinally ranked. We show that the unique equilibrium arising under this mechanism entails embellishment by both candidates and referees (a phenomenon known as the "Lake Wobegon effect"), but, despite this, always results in the selection of the most qualified candidate. Other notable features of the recruiting process, such as difficulty in distinguishing between candidates, changing perceptions of candidates over time, and the use of "wish lists" by recruiters are also accounted for.

JEL Classification Nos.: L20, J71, D82, J00

Keywords: Recruiting, Tournament, Embellishment

*The second author gratefully acknowledges financial support from the National Science Foundation. Address all correspondence to: John Moran, Center for Policy Research, Syracuse University, 426 Eggers Hall, Syracuse, NY 13244-1020. Email: jmoran@maxwell.syr.edu.

1 Introduction

Many organizations are faced with the problem of identifying the most qualified candidate to fill an available position. A crucial difficulty that recruiters face is that the underlying skills and abilities of the candidates are not directly observable; instead, employers must attempt to infer the abilities of candidates based on evidence provided by the candidates themselves. Such evidence might take the form of resumes, scores on standardized tests, or reports of past achievements. Much of the evidence that candidates provide is alterable (at some cost) in a manner that is undetectable by the recruiter. As a result, candidates have an incentive to engage in practices that "accentuate the positive" when reporting their qualifications, and recruiters must then decide to whom to offer the position on the basis of this possibly embellished information.

It is often argued that recruiters can solve this problem by soliciting recommendations from third-party "referees." However, those who are in the best position to provide information are often friends or close associates of the candidates, thus raising the prospect that referees, like the candidates themselves, may have incentives to overstate the qualifications of those they recommend.[1] Few job applicants, for example, receive a truly bad reference, while many enjoy recommendations which are highly favorable. Embellishment is also endemic in academic admission recommendations - so much so that standardized recommendation forms are increasingly designed to anticipate clustering in the upper percentiles of achievement. (Two such forms are illustrated in Figure 1; one from a large public university and another from a private university). This tendency toward overstatement has been termed the "Lake Wobegon effect," after the mythical town of the same name - a place where, "all the men are strong, all the women good looking, and all the children are above average."[2]

In light of this lack of veracity, recommendations play a surprisingly im-

[1] While few who provide recommendations admit to out-and-out lying, it seems clear that references often involve some degree of embellishment. When pressed, those who initially deny embellishing typically concede to having "emphasized the strong points" of the candidate. They also acknowledge their reluctance to volunteer unfavorable information, and admit to "downplaying" known weaknesses.

[2] Maxwell and Lopus (1994) use the term "Lake Wobegon effect" in conjunction with exaggerated claims of academic achievment on the part of college students. The town of Lake Wobegon was introduced by Garrison Keillor in his radio program, "A Prairie Home Companion" (see Keillor).

portant role in the choice among candidates. This is particularly true in the labor market, where roughly half of all workers find jobs through references provided by friends or relatives (Montgomery, 1991). Furthermore, workers who obtain jobs through personal referral earn higher initial salaries and have lower turnover than workers who find jobs through other means (Simon and Warner, 1992). These findings suggest that, despite the presence of the Lake Wobegon effect, third party evaluations do result in significant information transmission. This apparent contradiction, between the implied informational value of employee recruiting procedures on the one hand, and the systematic embellishment embodied in the Lake Wobegon effect on the other, raises two questions which from the focus of this paper. First, does the candidate recruitment process, plagued with misinformation and embellishment at all levels, lead to the selection of the most qualified candidates? And second, if embellishment is costly, and the recruiting procedure succeeds in selecting the most qualified candidate, why is the Lake Wobegon effect such a pervasive phenomenon?

To examine these questions we study a recruiting procedure commonly used by employers and other institutions. Specifically, we consider a mechanism consisting of an initial solicitation phase in which recommendations obtained from third-party referees are used to screen out all candidates who fall below a minimum ability threshold, followed by an interview phase in which the remaining candidates are ordinally ranked and the highest ranked candidate is hired. We show that this mechanism has a unique symmetric perfect Bayesian equilibrium which entails embellishment by both candidates and referees, yet always results in the selection of the most qualified candidate.

A key assumption of our model is that embellishment is costly. This assumption reflects a growing recognition in the literature that, in many settings, agents may need to expend resources to falsify the publicly revealed state of nature (see Lacker and Weinberg (1989), Maggi and Rodriguez-Clare (1995), and Crocker and Morgan, (1998)). Indeed, as Lacker and Weinberg have noted, "There are many instances in which 'lying' about the state of nature requires more than simply sending a false signal... Often, costly actions must be taken to lend credence to the signals being sent."[3] This seems especially true in a recruiting setting, where effective misrepresentation of a candidate's ability is likely to require greater effort than a truthful account-

[3] Lacker and Weinberg (1989), p. 1347.

ing. For example, in order to convincingly overstate a candidate's ability, a referee might need to expend additional effort to craft a strategically worded recommendation, or spend additional time lobbying on the candidate's behalf. Thus, when recommending two candidates of differing ability, a referee would have to expend greater effort on behalf of the less able candidate in order for the candidates to appear equally qualified.

In the same vein, the candidate himself may engage in costly activities in an effort to inflate a recruiter's perception of him. For the candidate, the nature of these activities will depend on the particular environment considered. When applying for a job, these might include time spent researching the employer, practicing interviewing skills, brushing up on buzzwords in relevant subjects, or polishing one's past employment or educational experiences, while in the case of academic admissions, test preparation services, such as those which prepare students for the SAT, MCAT, and LSAT, are a good example of a way in which candidates expend resources to appear more qualified than they actually are.[4] The distinguishing feature of each of these examples is that a candidate with lesser skill, ability, or knowledge will need to engage in more falsification in order to appear as qualified as a more able rival.

The mechanisms we study are, in effect, rank-order tournaments in that employers are able to distinguish, on a relative basis, the "paper" qualifications of candidates, but cannot directly observe their inherent "quality" or suitability for the position. Our results differ along several dimensions from standard tournament models, such as those of Lazear and Rosen (1981) and Rosen (1986).[5] Most obviously, our model is one of adverse selection rather than moral hazard. Thus the questions we consider focus on ensuring that the most suitable job candidate is selected, rather than inducing optimal effort on the part of a current employee. Second, attributes specific to the recruiting process lead to pooling outcomes in the first stage that are atypical of tournament models generally. The application of a multiple-stage tournament to employee selection is also, to our knowledge, novel.

[4] Some care in interpretation is required here. For our purposes, it is important that the activity in question not enhance the particular skills valued by the recruiter. Instead, we have in mind such things as last minute cramming and other short-term strategies that have no long-term effect on the candidate's underlying skill level, but which create the appearance that the candidate is more qualified than he is. Needless to say, this is not an issue when recommendations are provided by referees.

[5] For an excellent survey of the literature on tournaments, see McLaughlin (1988).

The paper proceeds as follows. In the next section we present a simple model of a widely observed two-stage recruiting mechanism. In section three we show that the unique symmetric equilibrium arising under this mechanism leads to optimal candidate selection despite the presence of embellishment by both candidates and referees (the Lake Wobegon effect). In addition, the model also offers a prediction about the pattern of embellishment observed in the screening stage of the game, and is shown to be consistent with a number of other notable features of the recruiting process, such as difficulty in distinguishing between candidates, changing perceptions of candidates over time, and the use of "wish lists" by recruiters. Section four contains concluding remarks. Longer proofs of key propositions are provided in Appendix A. In Appendix B, we examine candidate selection under four variations of the recruiting game: a one-stage procedure based on third-party recommendations alone; a one-stage procedure based on candidate interviews alone; a two-stage procedure which reverses the order in which candidates and referees are polled; and a two-stage procedure in which only the candidates themselves participate.

2 The Model

We consider the problem faced by an uninformed recruiter (employer, educational institution, etc.) who seeks to fill a position with the most qualified candidate.[6] There are N potential candidates, $i = 1, 2, ..., N$ for the position, each of whom brings to the market a given skill level $\theta \in [0, 1]$ (hereafter labeled "ability"), drawn from a known atomless distribution $F(\theta)$ with density $f(\theta) > 0$. The payoff to a recruiter when hiring a candidate of ability θ is $\pi(\theta)$, where $\pi(\cdot)$ is increasing in θ and $\pi(0) = 0$. A candidate who is hired receives a wage, w, that is set prior to the start of the recruiting process and remains fixed throughout.[7] We write the payoff function of the recruiter as

[6] While our model deals explicitly with the case in which there is a single vacancy to be filled, our results generalize to the case where there are l identical vacancies and each candidate can fill at most one vacancy.

[7] The assumption of a predetermined wage is consistent with observed practice across a variety of firms and industries. For instance, in fields such as accounting, banking and consulting, salaries are set in advance and remain largely beyond the control of those charged with making hiring decisions. The same is true in academic labor markets for junior faculty where, in many instances, salaries are determined separately from the process by which candidates are selected.

$\pi(\theta)$; which is net of w. In the event that the position is left vacant, the payoff to the recruiter is normalized to zero. There is a cost $k > 0$ to filling the position so the net benefit to a recruiter who hires a θ-type candidate is $\pi(\theta) - k$. If θ were known, the recruiter would only select a candidate of ability θ if $\pi(\theta) - k \geq 0$. Let $\theta_{min} \in (0,1)$ denote the minimum level of ability such that under full information the recruiter would be indifferent between filling the position or leaving it vacant. We will refer to candidates with ability less than θ_{min} as "unqualified."

In assessing the candidates, the recruiter might rely on information provided by a referee associated with candidate i, as well as from self-reporting by candidate i himself. Each candidate's ability is revealed to the candidate and the candidate's referee, but not to the recruiter or to the other candidates or referees. However, the distribution of candidate abilities is commonly known. At various stages in the hiring process candidates and referees may be asked to make reports (r) about the ability of the candidate. In doing so, we assume that both parties can, at some cost to themselves, effectively misrepresent the true ability of the candidate.[8] The cost to each party of falsifying the candidate's ability is given by the function $g(r - \theta)$. Our assumptions about the cost functions of candidates and referees are analogous to Lacker and Weinberg. Specifically,

A1. $$g(r - \theta) \begin{cases} = 0 & \text{if } r - \theta = 0 \\ > 0 & \text{if } r - \theta \neq 0 \end{cases}$$

A2. $$g'(r - \theta) \geq \gamma \text{ for } r - \theta \geq 0 \text{ where } \gamma > 0$$

A3. $$g''(\cdot) > 0$$

(A1) implies that misrepresenting the candidate's ability is more costly than reporting truthfully. (A2) and (A3) imply that falsification costs are increasing in the size of the misrepresentation, and that g is invertible over the domain $r - \theta \geq 0$. It is worth noting that while the assumption $g'(0) > 0$ is shared with the model of Lacker and Weinberg (1989), in their model this condition is needed to generate regions of no-falsification in the contracts

[8] Specific falsification activities are discussed in the introduction.

they consider. In our model, this assumption is a technical one used to ensure that equation (2) (shown below) is well-behaved at the boundary. As a subsequent example highlights, the first-stage game results are not affected by weakening this assumption to $g'(0) \geq 0$:

To highlight the effects of falsification on information transmission without confounding insurance or risk considerations, we assume that all candidates and referees are risk neutral.[9] Thus, the utility of candidate i of type θ who submits a report r is

$$U_i(r;\theta) = \begin{cases} w - g(r-\theta) & \text{if hired} \\ -g(r-\theta) & \text{if not hired} \end{cases}$$

where w is the reward (wage) in units of a numeraire good if i obtains the position and g is the falsification cost function denominated in units of the numeraire good. Similarly the utility of i's associated referee is

$$V_i(r;\theta) = \begin{cases} v - g(r-\theta) & \text{if } i \text{ is hired} \\ -g(r-\theta) & \text{if } i \text{ is not hired} \end{cases}$$

where v is the reward in units of the numeraire good to i's referee if i obtains the position and g is the falsification cost function denominated in units of the numeraire good. In specifying the referee's utility function, we have assumed that the referee receives some benefit from having their candidate obtain the position. This benefit may either be tangible, as would be the case if the candidate were able to use their position to provide assistance to the referee, or intangible, as would occur if the referee's preferences reflect a degree of altruism towards the candidate.

We assume that the recruiter commits to the mechanism used to select candidates and study the equilibrium properties of a class of mechanisms frequently observed in practice. Specifically, the recruiting mechanism consists of two stages: an initial stage in which unqualified candidates are removed from consideration; and a second stage in which the remaining candidates

[9] Since there are two discrete outcomes that affect a candidate's utility (being hired vs. not being hired), our model readily extends to cases where candidates have preferences of the form:

$$U_i(r;\theta) = u(w) - g(r-\theta)$$

where u is a strictly concave function. To see this, notice that we may normalize $u(0) = 0$ and $u(w) = 1$ for preferences in this class. An identical argument holds for the preferences of referees. This, in effect, yields the preferences we consider.

are ordinally ranked. The highest ranked candidate in the second stage obtains the position. In the event of a tie in the second stage, one of the tied candidates is selected at random.

For the remainder of the paper, we will assume that the first stage of the process is based on third-party recommendations while the second stage entails direct interviewing of the candidates. In Appendix B, we demonstrate that our results do not depend on the order in which candidates and referees are polled, however, both stages of the process are required to ensure optimal candidate selection.

3 Equilibrium Characterization

3.1 Interviews

We begin with the second-stage (interview) game. Suppose that in the recommendation stage of the process all candidates whose abilities were below some threshold θ' were eliminated from consideration. Each of the remaining candidates competes for reward w associated with obtaining the position. Candidates simultaneously make reports about their ability and the cost to a candidate of type θ of making a report r is $g(r - \theta)$. The position is awarded to the candidate who reports the highest ability, which can be interpreted as having appeared most qualified during the interview.

3.1.1 Equilibrium

Since the game is symmetric, we restrict attention to symmetric equilibrium reporting strategies. We begin by showing that all symmetric equilibrium reporting strategies are monotonic.

Proposition 1 *Suppose $\rho(\theta)$ is a symmetric equilibrium reporting strategy to the second-stage game. Then for all θ, $\rho(\theta)$ is non-decreasing.*

Proof. See Appendix A.

Next, suppose without loss of generality that when candidates $2, 3, ..., N$ follow the reporting strategy $\rho(\theta)$, candidate 1 chooses a report r to maximize:

$$U(r|\theta) = wF\left(\rho^{-1}(r)\right)^{N-1} - g(r - \theta)$$

where $F(\rho^{-1}(r))^{N-1}$ is the probability of being hired given a report of r.

Differentiating we obtain

$$\frac{w(N-1)f(\rho^{-1}(r))F(\rho^{-1}(r))^{N-2}}{\rho'(\rho^{-1}(r))} - g'(r-\theta) = 0$$

Using symmetry,

$$\frac{w(N-1)f(\theta)F(\theta)^{N-2}}{\rho'(\theta)} - g'(\rho(\theta)-\theta) = 0$$

Rearranging, we obtain the differential equation:

$$\rho'(\theta) = \frac{wh_{N-1}(\theta)}{g'(\rho(\theta)-\theta)} \qquad (1)$$

where $h_{N-1}(\theta)$ is the density of the highest of $N-1$ independent draws from $F(\theta)$.

Since equation (1) is a first-order nonlinear differential equation, it is useful to define $\psi(r;\theta) \equiv \frac{wh_{N-1}(\theta)}{g'(r-\theta)}$.

We make the following technical assumptions, which do not affect the economics of the results, but considerably simplify the characterization of an equilibrium.

A 4. ψ is Lipschitz.

A 5. There exists a $\bar{\theta} \in (0;1)$ such that for $\theta \leq \bar{\theta}$; $wh_{N-1}(\theta) \leq \gamma$ and for $\theta > \bar{\theta}$; $N\gamma > wh_{N-1}(\theta) > \gamma$.

The equilibrium reporting strategy is characterized in Proposition 2.

Proposition 2 The unique symmetric Bayesian Nash equilibrium of the second-stage game consists of a reporting strategy, $\rho(\theta)$, where:

$$\rho(\theta) = \theta \qquad \text{for } \theta < \bar{\theta}$$

and $\rho(\theta)$ solves the differential equation

$$\rho'(\theta) = \frac{wh_{N-1}(\theta)}{g'(\rho(\theta)-\theta)} \qquad (2)$$

and endpoint condition:[10]

$$\begin{matrix} \rho\left(\hat{\theta}^1\right) = \bar{\theta} & \text{if } \hat{\theta}^1 \geq \theta^0 \\ \rho(\theta^0) = \theta^0 & \text{if } \hat{\theta}^1 < \theta^0 \end{matrix} \qquad \text{for } \theta \geq \bar{\theta}$$

Proof. See Appendix A.

Proposition 2 implies, among other things, that candidates of sufficiently low ability will not find it in their interest to embellish their (modest) qualifications. In particular, when $\hat{\theta}^1 > \theta^0$; candidates with abilities from θ^0 to $\hat{\theta}^1$ will choose to report truthfully.[11]

In Proposition 3, we establish that embellishment is an integral part of any symmetric equilibrium. Indeed, when $\theta^0 > \hat{\theta}^1$; almost all candidates choose to embellish.

Proposition 3 *Embellishment occurs for almost all candidates with ability* $\theta \in \left(\max\left(\bar{\theta};\theta^0\right);1\right]$:

Proof. For all $\theta > \max\left(\bar{\theta};\theta^0\right)$; if $\rho(\theta) = \theta$; then $\rho'(\theta) > 1$: Hence, there can be only finitely many points where embellishment does not occur. ■

Some properties of the equilibrium strategies are worth noting. First, since reporting strategies are monotonic, it follows that the ordinal ranking obtained from the candidates' embellished qualifications will be identical to

[10] Recall that candidates with $\theta < \theta^0$ were removed from consideration during the first round. Thus, if $\theta^0 \geq \bar{\theta}$, then it is automatic that $\theta \geq \bar{\theta}$.

[11] An alternative interpretation of the second-stage competition is as follows: Suppose that n bidders compete in a single-object, first-price all-pay auction where it is common knowledge that the object is worth w to each bidder. Bidders differ in their skill (θ) in preparing bids in the following fashion: A bidder may, at no cost, make a bid equal to his type, but bids above one's type are increasingly costly. Since submitting a bid of a given level is more expensive for less skilled bidders, bid strategies that are increasing in a bidder's type arise in equilibrium. Thus, efficient sorting (in the form of allocating the object to the bidder with the highest type) is achieved by virtue of the fact that "signaling"- in the form of bidding to a certain level - is differentially costly across types. At first glance, this may appear to be different from many applications in contract theory in which a type-dependent transfer payment is required in order to affect separation. Here, the transfer is not type-dependent per se, but it is probabilistically type-dependent due to the fact that the candidate's likelihood of obtaining the position is an increasing function of θ:

the ordinal ranking of the candidates' true abilities. The underlying economics of this follows from the fact that since marginal embellishment costs are increasing, a level of publicly observable credentials that is profitable for a lower ability type will be even more profitable for a higher ability candidate. To see this, notice that to generate the same level of publicly revealed qualifications, a lower type must expend greater resources than a higher ability candidate, while the expected benefit of presenting identical public qualifications are exactly the same. Thus, competition for the position ensures that the publicly observed credentials of lower ability candidates will never outshine those with higher ability. Furthermore, the form that this competition takes implies that a positive measure of candidates will always choose to "accentuate the positive" and report higher than their true ability. In other words, on average, candidates are above average, a finding which is consistent with the Lake Wobegon effect.

To obtain some additional intuition for why Lake Wobegon effects necessarily arise in this context, it is instructive to examine why truth-telling is not an equilibrium. Suppose that all of the candidates were initially following a truth-telling strategy. Now, a candidate with sufficiently low ability would perceive the gains from a small amount of embellishment to be fairly low (since it is likely that his ability is significantly below that of the highest ability applicant). Since the marginal costs of even a small amount of falsification are positive, these low ability candidates would be content with truth-telling. A candidate with higher (but not the highest) ability would be more optimistic about his prospects of "embellishing his way to the top," and hence, would be willing to incur the costs of engaging in some falsification. Obviously, candidates in the upper tail of the ability distribution would perceive little upside in falsifying, and hence, would be content to report honestly. Thus, it is the candidates of medium ability who have the strongest incentives to embellish. In equilibrium, higher ability candidates recognize these incentives and react by falsifying their own types, since by telling the truth they would be viewed as medium types in the eyes of the recruiter and would therefore face the same incentives to falsify as the medium ability candidates. This effect "reverberates" all the way to the highest ability candidates who must then embellish simply to distinguish themselves from medium types. The upshot is that the resulting equilibrium preserves the ordinal ranking of candidates while at the same time leading to systematic embellishment (the Lake Wobegon effect) on the part of candidates.

Finally, observe that the reporting strategies arising in the interview stage do not require recruiters to "invert back" or in any way decipher the reports issued by the candidates. Instead, recruiters are free to delegate interviewing duties to unsophisticated agents who will simply hire the candidate who appears most qualified during the interview.

To see all of this more easily, consider the following example. Suppose that there are 4 candidates competing for a position whose value is 1. The candidates' abilities are i.i.d. draws from the uniform distribution on $[0;1]$. Candidates have been pre-screened so that only candidates whose reported ability exceeds $\frac{1}{2}$ are being interviewed. Finally, let $g(r-\theta) = \frac{3}{4}(r-\theta)^2 + (r-\theta)$ for $r \geq \theta$. One may readily verify that Φ is Lipschitz and that a $\bar{\theta}$ satisfying (A5) occurs at $\theta = \frac{1}{2}$. The equilibrium strategy $\rho(\theta)$ is implicitly characterized by:

$$\rho' \times \left(\frac{1}{2}(\rho - \theta) + \frac{1}{4} \right) = \theta^2$$

$$\rho\left(\frac{1}{2}\right) = \frac{1}{2}$$

A numerical solution to this system of equations (illustrated in Figure 2) highlights the discrepancy between truth-telling, represented by the thin line, and the qualification levels claimed in equilibrium, represented by the thick line. Note that despite the existence of systematic embellishment, the relative ranking of candidates is preserved in equilibrium. Thus, the recruiter can obtain full ordinal revelation and hire the highest ability candidate in the second-stage game, even in the presence of the Lake Wobegon effect.

3.2 Recommendations

Now consider the first-stage game. The recruiter's objective is to choose the recommendation threshold r^* for advancing candidates to the second stage in a manner which maximizes expected payoffs.

3.2.1 Equilibrium

To remove unqualified candidates from further consideration (i.e. those with $\theta < \theta_{min}$), the recruiter seeks to create a short list of candidates with ability above some threshold $\mu^0 \geq \mu_{min}$.[12] To achieve this, the recommendation threshold r^* must be set in a manner which leaves referees just indifferent between embellishing to meet the threshold for the marginal types. The indifference condition for referees of the marginal types is given by

$$
\begin{aligned}
vH_{N-1}(\mu^0) - g(r^* - \mu^0) &= 0 \\
r^* &= g^{-1}(vH_{N-1}(\mu^0)) + \mu^0
\end{aligned}
$$

where $H_{N-1}(\mu^0)$ is the distribution function of the highest of $N-1$ independent draws from $F(\mu)$. Observe that the chosen recommendation threshold r^* is always above the underlying ability threshold μ^0 used to advance candidates to the second round. Thus, in the presence of the Lake Wobegon effect, recruiters find it in their interest to specify a minimum reported qualification level which exceeds the actual qualifications (ability) needed to adequately function in the job. This appears to be a common strategy in practice, as employers frequently enumerate "wish lists" of job qualifications in excess of what is ultimately deemed acceptable.

It is immediate that this cutoff rule implies that equilibrium recommendation strategies in the first-stage game are:

Proposition 4 Referees submit recommendations according to:

$$
\begin{aligned}
r(\mu) &= \mu \quad \text{if } \mu \in [0; \mu^0) \\
&= r^* \quad \text{if } \mu \in \left[\mu^0; g^{-1}(vH_{N-1}(\mu^0)) + \mu^0\right] \\
&= \mu \quad \text{if } \mu \in (g^{-1}(vH_{N-1}(\mu^0)) + \mu^0; \infty]
\end{aligned}
$$

and believe that the second-stage game will be resolved according to Proposition 2.

To summarize, the use of a cutoff rule by the recruiter results in recommendations that are increasing in the true ability of the candidate but which exhibit pooling at some cutoff level r^*. Lower ability candidates are honestly reported, as are "superstars" whose true ability lies above the threshold r^*. The middle range of candidates, whose ability lies between μ^0 and r^*, are all uniformly reported as being of ability r^*. Thus, the average ability reported by referees is greater than the true statistical average ability of the candidates, leading to a Lake Wobegon effect in the recommendation stage

[12] For now, we allow μ^0 to remain arbitary; later, we will set it to its optimal level, μ_{min}.

of the game as well. Moreover, the existence of a mass point of candidates, all of whom receive recommendation r^*, means that a subset of candidates of differing abilities will initially appear indistinguishable to the recruiter. This is in keeping with the often-expressed view that candidates can be difficult to distinguish based on their recommendations alone.

For optimal candidate selection, the recruiter needs to specify a threshold that screens out all candidates whose abilities are below $\theta_{min} = \pi^{-1}(k)$. Thus, $\theta' = \theta_{min}$, and the optimal threshold is:

$$r^* = g^{-1}\left(vH_{N-1}\left(\pi^{-1}(k)\right)\right) + \pi^{-1}(k).$$

This expression leads to the obvious implications that screening standards will become more stringent when either the cost of filling the position increases, or the distribution of ability shifts in a manner consistent with first-order stochastic dominance.

Although we have assumed that falsification costs satisfy A1-A3, pooling by referees will occur even if we relax these assumptions.[13] The following example is illustrative. Suppose that there are four candidates each of whose ability is an i.i.d. draw from the uniform distribution on $[0,1]$. The referees obtain utility of 2 by successfully placing their candidates and have falsification costs $g = (r - \theta)^2$.

Qualified candidates are those whose ability exceeds $\frac{1}{2}$. This induces a recommendation threshold of

$$r^* = 1$$

Thus, the average reported ability of candidates is

$$E(r) = \frac{5}{8}$$

as opposed to a true average of $\frac{1}{2}$. Finally, half the candidates receive recommendations indicating that they have the highest conceivable level of ability; thus, on paper, candidates will be almost "too good to be true."

Notice that the bunching of candidate types which occurs during the recommendation stage gets sorted out in the interview stage. That is, the lower ability candidates tend to perform worse in the interview than would

[13] Specifically, the assumption that $g'(0) > 0$ is made solely to generate well behaved differential equations characterizing reporting in the second-stage game. This example shows that for the first-stage game, this assumption may be dropped.

have been envisioned based on their recommendations. In figure 2, these candidates are those whose abilities lie between .5 and .825. At the same time, the higher ability candidates will generally exceed the expectations created by their recommendations. As a result, the model is consistent with a phenomenon whereby candidates "come across" differently in the interview and recommendation phases of the process, leading recruiters to change their rankings in the aftermath of the interviews.

4 Conclusion

The analysis presented here demonstrates that a widely observed two-stage recruiting procedure can be used to identify the most qualified candidate for a position, even when the candidates' qualifications are subject to embellishment by both the candidates themselves and third parties chosen to evaluate them. Optimal candidate selection arises as a consequence of the differential costs of embellishing the qualifications of candidates with different underlying abilities. This induces a sorting among the embellished qualifications of candidates that is ordinally revealing. That is, if falsification costs rise with the degree of embellishment, competition among candidates leads to an ordinal ranking of embellished qualification levels that is identical to the ordinal ranking of the candidates' true qualifications. An implication is that recruiters who choose the best candidate "on paper" indeed obtain the most qualified individual. Given this decision rule by recruiters, a candidate who refrained from embellishment would be leapfrogged in the selection process by less qualified candidates; thus, the Lake Wobegon effect can be seen as a consequence of the negative externalities associated with competition for rents from the desired employment opportunity.

In the first stage of the hiring process, the model predicts a clustering of the recommendations made on behalf of candidates with average abilities. This seems consistent with the view that candidates of different abilities are often initially hard to distinguish. The model also provides a rationale for the practice of specifying a "wish list" of minimum job qualifications in excess of those actually needed for adequate performance in the job; that is, since the recruiter anticipates receiving embellished reports on behalf of average candidates, it is essential to raise the stated qualification threshold above the actual threshold needed to ensure adequate performance. Finally, because the partial pooling of candidates that occurs in the first stage is resolved in the second stage, the model offers an explanation for why recruiters' rankings of candidates are likely to change as the process unfolds.

The model presented here is, of course, a highly stylized version of employee recruitment. Thus, it is useful to consider the robustness of our results to perturbations and extensions of the model. Throughout, we have made the simplifying assumption that the application of falsification effort leads to a deterministic outcome in terms of publicly revealed ability. We can easily relax this assumption to allow for a stochastic idiosyncratic component to revealed ability that is independent of falsification activities. Apart from endpoint effects, such an extension will have no impact on marginal falsification incentives and hence, subject to the noise component of ability revelation, all of the above results will continue to hold. Note, however, that such a noise term would not allow for beliefs that perfectly "invert back" in a signaling version of the model, thus highlighting the difference between our tournament-like approach and one based on signaling.

Another extension might be to have referees and candidates receive imperfectly correlated signals of the candidate's true ability. In this case, recruiters would have to consider not only the cost of utilizing the "interview" technology, but also the relative precision of the signals held by the two parties. Nonetheless, both the Lake Wobegon effect and our finding that candidates are accurately ranked in the second stage of the process would still be present in a symmetric equilibrium.

Finally, while the mechanism we study always leads to the selection of the most qualified candidate, it also entails wasteful falsification activities on the part of both candidates and referees, and requires that the short list of candidates be evaluated twice (first, by third-party referees, and later, via direct interviewing). Thus, an interesting question is what constitutes an optimal recruiting mechanism when one is cognizant of the costs to both candidates and recruiters of administering the recruiting procedure? Answering this question is beyond the scope of the present paper and remains for future research.

References

[1] Crocker, Keith and John Morgan, 1998. "Is Honesty the Best Policy? Curtailing Insurance Fraud through Optimal Incentive Contracts," Journal of Political Economy, 106, 355-375.

[2] Hirsch, Morris and Stephen Smale, 1974. Di¤erential Equations, Dynamical Systems, and Linear Algebra (Academic Press, New York).

[3] Keillor, Garrison. "Prairie Home Companion," signature line, Minnesota Public Radio.

[4] Lacker, Je¤rey and John Weinberg, 1989."Optimal Contracts under Costly State Falsi…cation," Journal of Political Economy, 97, 1345-1363
.

[5] Lazear, Edward and Sherwin Rosen, 1981. "Rank-Order Tournaments as Optimum Labor Contracts," Journal of Political Economy, 89, 841-864 .

[6] Maggi, Giovanni and Andres Rodriguez-Clare, 1995. "Costly Distortion of Information in Agency Problems," Rand Journal of Economics, 26, 675-689.

[7] Maxwell, Nan and Jane Lopus, 1994. "The Lake Wobegon E¤ect in Student Self-Reported Data," American Economic Review - Papers and Proceedings, 84, 201-205.

[8] McLaughlin, Kenneth, 1988, "Aspects of Tournaments Models: A Survey," In: Ronald Ehrenberg, (Eds.), Research in Labor Economics, Volume 9. JAI Press, Greenwich, Connecticut and London, 225-256.

[9] Milgrom, Paul and Robert Weber, 1982. "A Theory of Auctions and Competitive Bidding," Econometrica, 50, 1089-1122.

[10] Montgomery, James, 1991. "Social Networks and Labor-Market Outcomes: Toward an Economic Analysis," American Economic Review, 81, 1407-1418.

[11] Rainville, Earl and Phillip Bedient, 1981, Elementary Di¤erential Equations, Sixth Edition, (Macmillan, New York).

[12] Rosen, Sherwin, 1986. "Prizes and Incentives in Elimination Tournaments," American Economic Review, 76, 701-715.

[13] Simon, Curtis and John Warner, 1992. "Matchmaker, Matchmaker: The Effect of Old Boy Networks on Job Match Quality, Earnings, and Tenure," Journal of Labor Economics, 10, 306-330.

[14] Spence, Michael, 1973. "Job Market Signaling," Quarterly Journal of Economics, 87, 355-374.

5 Appendix A - Proofs

5.1 Proof of Proposition 1

Suppose that the contrary is true. First, notice that strategies such that $\rho(\theta) < \theta$ are not rationalizable. Hence, there exists $\theta, \widehat{\theta}$ with associated strategies $\rho(\theta)$ and $\rho\left(\widehat{\theta}\right)$ such that $\theta < \widehat{\theta}$, $\rho(\theta) > \rho\left(\widehat{\theta}\right) \geq \widehat{\theta}$ and $\left|\theta - \widehat{\theta}\right| < \varepsilon$ for small ε.

By convexity of the g function, we have that

$$g(\rho(\theta) - \theta) - g\left(\rho\left(\widehat{\theta}\right) - \theta\right) > g\left(\rho(\theta) - \widehat{\theta}\right) - g\left(\rho\left(\widehat{\theta}\right) - \widehat{\theta}\right).$$

Incentive compatibility requires that

$$F(\theta)^{N-1}v - g(\rho(\theta) - \theta) \geq F\left(\widehat{\theta}\right)^{N-1}v - g\left(\rho\left(\widehat{\theta}\right) - \theta\right)$$

and

$$F\left(\widehat{\theta}\right)^{N-1}v - g\left(\rho\left(\widehat{\theta}\right) - \widehat{\theta}\right) \geq F(\theta)^{N-1}v - g\left(\rho(\theta) - \widehat{\theta}\right)$$

Hence,

$$g\left(\rho(\theta) - \widehat{\theta}\right) - g\left(\rho\left(\widehat{\theta}\right) - \widehat{\theta}\right) \geq g(\rho(\theta) - \theta) - g\left(\rho\left(\widehat{\theta}\right) - \theta\right)$$

but this is a contradiction and playing $\rho(\theta)$ for type $\widehat{\theta}$ is a profitable deviation.

5.2 Proof of Proposition 2

First, note that since all candidate types $\theta < \theta^0$ are eliminated from consideration in the second stage, it is sufficient to characterize bidding functions for types in the interval $[\theta^0, 1]$. We begin by proving several lemmas which place restrictions on the characteristics of equilibrium reporting strategies.

Lemma 5 No equilibrium reporting strategy can contain atoms.

Proof. Suppose not. Then there exists some open interval of types, (θ_1, θ_2), that make the same report, r, in equilibrium. Suppose that the highest ability candidate making report r is θ_2 and the lowest is θ_1. In this case a candidate of ability $\theta \in (\theta_1, \theta_2)$ obtains expected payoffs that are strictly less than $vF(\theta_2)^{N-1} - g(r - \theta)$. However, by sending a slightly higher report r', such a candidate can earn arbitrarily close to $vF(\theta_2)^{N-1} - g(r-\theta)$. This is a profitable deviation. ■

Since equilibrium reporting strategies contain no atoms and are increasing, we may then deduce that equilibrium reporting strategies are strictly increasing.

Lemma 6 For all candidates of ability $\theta \in (\theta^0, 1)$, any equilibrium reporting strategy is continuous.

Proof. Suppose not, then there exists a type θ_1 such that $\lim_{\theta \uparrow \theta_1} \rho(\theta) < \lim_{\theta \downarrow \theta_1} \rho(\theta)$. If $\rho(\theta_1) > \lim_{\theta \uparrow \theta_1} \rho(\theta)$, then a candidate of ability θ_1 can lower his costs by a finite amount by choosing the report $r = \lim_{\theta \uparrow \theta_1} \rho(\theta)$, without affecting his probability of winning. If $\rho(\theta_1) < \lim_{\theta \downarrow \theta_1} \rho(\theta)$, then candidates with abilities slightly above θ_1 can, by reducing their reports to $r = \rho(\theta_1)$, save a finite amount in falsification costs while reducing their chances of winning by an arbitrarily small amount. Both of these are profitable deviations. ■

Since reporting strategies are strictly increasing and continuous, they must be almost everywhere differentiable.

Lemma 7 Equilibrium reporting strategies must have the property that $\rho(\theta^0) = \theta^0$.

Proof. To see this, suppose that an equilibrium strategy had the property that $\rho(\theta^0) > \theta^0$. In that case, since there are no atoms, a candidate can lower his costs by reporting $\rho(\theta^0) = \theta^0$ without affecting his probability of winning. This is a profitable deviation. ■

Given lemmas 5, 6, and 7, we may now proceed to the proof of Proposition 2. First, note that since Φ is Lipschitz, a solution to (2) exists (Rainville and Bedient (1981), p. 298). Moreover, since $\rho'(\theta) \geq 0$ for all θ, such a solution is increasing. By Theorem 1 in Chapter 15 of Hirsch and Smale, we know that the solution to the differential equation is locally unique. By standard techniques the local solution may be extended (see Hirsch and Smale, chapter 8, Section 5).

Next we show that such a solution is a BNE. Suppose that candidate 1 pretends that his type is z when his true ability is θ and all other candidates are employing the strategy ρ. Clearly, any strategy $\rho(z) < \theta$ is dominated by reporting truthfully; hence we restrict attention to cases where $\rho(z) \geq \theta$. Candidate 1's marginal utility is

$$MU_1(z|\theta) = \frac{w h_{N-1}(z)}{\rho'(z)} - g'(\rho(z) - \theta)$$

Evaluating the first-order condition at $\theta = z$ and substituting in for $\rho'(z)$ yields

$$MU_1(z|\theta) = g'(\rho(z) - z) - g'(\rho(z) - \theta)$$

Recall that $\rho(z) \geq z$; hence, for $z < \theta$, $MU_1(z|\theta) > 0$ and for $z > \theta$, $MU_1(z|\theta) < 0$. Thus, ρ is a BNE.[14]

Since any symmetric equilibrium reporting strategy is almost everywhere differentiable and has the endpoint condition $\rho(\theta^0) = \theta^0$, we know that equation (2) is necessary. Thus, we may apply standard differential equation uniqueness theorems (see above) to establish that $\rho(\cdot)$ is the unique symmetric increasing equilibrium. ■

6 Appendix B - Alternative Extensive Forms

1. Suppose that in the first stage candidates are polled, followed by a short-list of referees. Then, the properties of the mechanism are identical. To see this, notice that the only difference between the utilities of referees versus candidates was in their utility from obtaining the position (or from supporting a candidate who obtained the position). In short, if we changed the extensive form of the game to interviews followed by recommendations, and adopted a cutoff strategy in the interview phase, all of the preceding results would continue to hold.

2. Suppose that only the second-stage procedure took place. Then, the highest ability candidate would still be selected since Propositions 1 and 2 hold for the case where the cutoff in the first-stage game was $r^* = 0$. However, the ability of this candidate need not exceed θ_{min}; hence, we do not obtain optimal candidate selection with this procedure.

[14] Relaxing A5 results in the possibility of: (1) multiple intersections with the $\rho(\theta) = \theta$ boundary; or (2) truth-telling everywhere (if w is small relative to γ). Possibility (1) simply introduces additional technical complexity without altering the qualitative ordinal revelation result, while possibility (2) represents the economically uninteresting case where falsification is too expensive to employ.

3. Suppose that only the first-stage procedure took place and a candidate was selected randomly from among those whose recommendations exceeded r^*. In this case, the indifference condition for referees is

$$\frac{v}{(n-1)(1-F(\theta'))} - g(r^* - \theta') = 0.$$

and θ' solves

$$\frac{v}{(n-1)} = g(r^* - \theta')(1 - F(\theta')).$$

For $\theta < r^*$, $g(r^* - \theta)(1 - F(\theta))$ is decreasing, and for $\theta = r^*$, $g(r^* - \theta)(1 - F(\theta)) = 0$. Hence, provided $g(r^*) > \frac{v}{(n-1)}$, there exists a unique $\theta' \in (0, r^*)$ solving the above equation and we can write the recruiter's maximization solely as a function of θ'. Thus, the recruiter selects θ' to maximize

$$\begin{aligned}
&(E(\pi(\theta) \mid \theta > \theta') - k)\Pr(\max(\theta_1, \theta_2, \ldots, \theta_n) > \theta') \\
&= \left(\frac{1}{1-F(\theta')}\int_{\theta'}^{1} \pi(t) f(t)\,dt - k\right)\left(1 - F(\theta')^n\right).
\end{aligned}$$

Differentiating with respect to θ', a necessary condition is

$$\begin{aligned}
&\left(1 - F(\theta')^n\right)\left(\frac{f(\theta')}{(1-F(\theta'))^2}\int_{\theta'}^{1} \pi(t) f(t)\,dt - \frac{\pi(\theta') f(\theta')}{1-F(\theta')}\right) \\
&- nF(\theta')^{n-1} f(\theta')\left(\frac{1}{1-F(\theta')}\int_{\theta'}^{1} \pi(t) f(t)\,dt - k\right) \\
&= 0.
\end{aligned}$$

and θ' solves

$$\frac{\left(1 - F(\theta')^n\right)}{\left(nF(\theta')^{n-1}\right)(1 - F(\theta'))} = \frac{E(\pi(\theta) \mid \theta > \theta') - k}{E(\pi(\theta) \mid \theta > \theta') - \pi(\theta')}.$$

It is immediate that $\theta' > \theta_{min}$; hence, for a given cost of filling the position, the stringency of standards in the first-stage procedure is higher than under the two stage mechanism. It is also immediate that optimal candidate selection does not occur with such a procedure.

4. Consider the case where the candidates themselves undergo the two-stage procedure. One may think of this as candidates submitting resumes for the position in the first stage. Candidates whose resumes exceed some

threshold are then invited for face-to-face interviews. Suppose that the first-stage procedure screens all candidates with ability below θ' in equilibrium. Clearly, second-stage reporting strategies are identical to those in Proposition 2. Thus, the expected surplus of a candidate with ability θ in the second stage is

$$S(\theta) = wF(\theta)^{N-1} - g(\rho(\theta) - \theta)$$

where $\rho(\theta)$ is the reporting strategy given in Proposition 2. When $\theta \leq \theta^1$; $\rho(\theta) = \theta$ so it follows that $S(\theta)$ is increasing in θ. When $\theta > \theta^1$; equation (2) holds and, upon substitution, it follows that

$$\begin{aligned} S'(\theta) &= g'(\rho(\theta) - \theta) \\ &> 0 \end{aligned}$$

Thus, the expected surplus is increasing in ability.

Suppose that in the first stage, a threshold r^* is employed. If a candidate of type θ' embellishes up to r^* (or greater) in the first stage, she anticipates receiving surplus of $S(\theta')$ in the second stage. The cost of embellishing up to r^* in the first stage is $g(r^* - \theta')$. Thus, for a candidate of type θ' to be indifferent between embellishing up to r^* in the first stage and not embellishing to this level requires that

$$S(\theta') - g(r^* - \theta') = 0.$$

Obviously the cost of embellishing up to r^* is greater for candidates with ability below θ'; and, as we showed above, surplus is increasing in θ. Thus, no candidate with ability below θ' will choose to embellish up to r^*. A similar argument implies that all candidates with abilities above θ' will choose to embellish to at least that level in the first stage. As a result, we have shown that in a two-stage procedure where candidates themselves participate at both stages, an optimal sorting still holds. Indeed, notice that when $v = w$; the threshold to induce optimal sorting is set at exactly the same level as when the referee and the candidate are separate players. This is because the marginal type does not falsify in the second-stage interviews, so $S(\theta') = wF(\theta)^{N-1} = vH_{N-1}(\theta')$. In general, the optimal threshold will depend on w.

Figure 1 – Clustering in recommendations

2. Evaluation: In comparison with other students in the same field who have the same amount of experience and training, I rate this person as follows:

	Top 5%	Top 10%	Top 20%	Upper 50%	Unable to rate
Knowledge in proposed subject of study					
Ability to grasp new concepts					
Originality, intellectual creativity					
Mathematical and logical thought					
Written expression					
Oral expression					
Laboratory skills (if applicable)					
Perseverance towards goals					
Potential as teacher (if applicable)					

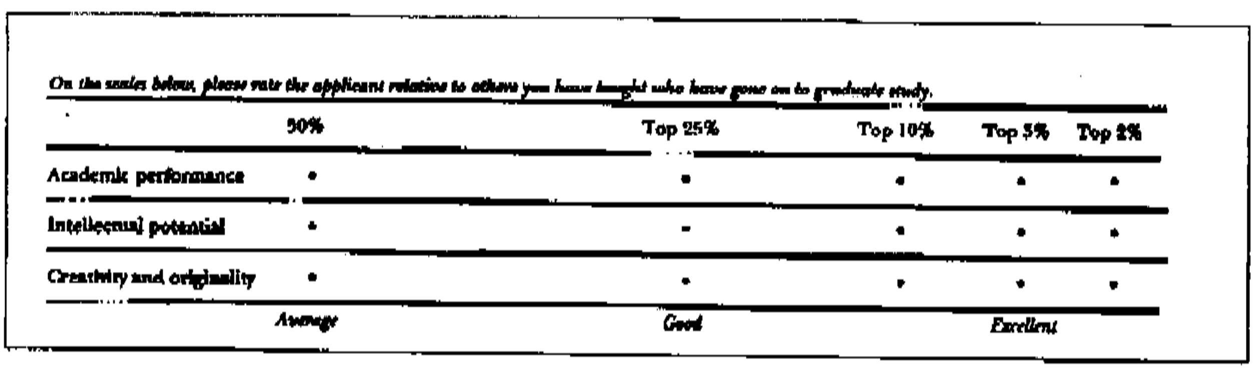

On the scales below, please rate the applicant relative to others you have taught who have gone on to graduate study.

	50%	Top 25%	Top 10%	Top 5%	Top 2%
Academic performance	•	•	•	•	•
Intellectual potential	•	•	•	•	•
Creativity and originality	•	•	•	•	•
	Average	*Good*		*Excellent*	

Figure 2 – Equilibrium Lake Wobegon Effects