

Inflated Applicants: Attribution Errors in Performance Evaluation by Professionals

Samuel A. Swift¹

Don A. Moore²

Zachariah S. Sharek¹

Francesca Gino³

¹*Carnegie Mellon University*, ²*University of California at Berkeley*, ³*Harvard University*

The authors appreciate the support of National Science Foundation Grant SES-045173, a doctoral student fellowship from the Graduate Management Admissions Council, and a SURG research grant from Carnegie Mellon University. Address correspondence to samswift@alumni.cmu.edu.

Abstract

When explaining others' behaviors, achievements, and failures, it is common for people to attribute too much influence to disposition and too little influence to structural and situational factors. We examine whether this tendency leads even experienced professionals to make systematic mistakes in their selection decisions, favoring alumni from academic institutions with lenient grading and employees from forgiving business environments. We find that candidates benefiting from favorable situations are more likely to be admitted and promoted than their equivalently skilled peers. The results suggest that decision-makers take high nominal performance as evidence of high ability and do not discount it by the ease with which it was achieved. These results clarify our understanding of the correspondence bias using evidence from both archival studies and experiments with experienced professionals. There are profound implications for both admissions and personnel selection practices.

Keywords: correspondence bias; selection decisions; evaluations; attribution

Inflated Applicants: Attribution Errors in Performance Evaluation by Professionals

The National Association of Colleges and Employers reports that 66 percent of employers screen candidates by grade point average (GPA), and 58 percent of employers indicated that a GPA below 3.0 all but eliminates a candidate's chances of being hired (Koeppel, 2006). Academic grades, after all, test the mental ability central to predicting job performance (Schmidt & Hunter, 1998). However, by setting a cutoff GPA for all candidates, employers implicitly assume that a grade from one school is equivalent to a grade from another. This is a problematic assumption because universities vary considerably in their grading standards; otherwise similar universities display large differences in the grades they award their students (Attewell, 2001; Bagues, Labini, & Zinovyeva, 2008; Goldman & Widawski, 1976; Rojstaczer & Healy, 2010; Strenta & Elliott, 1987). Accounting for variation in grading leniency between institutions represents a special case of a more general problem: how to make attributions about performance while considering both the performer's abilities and the difficulty of the task.

Organizations encounter this vexing attribution problem routinely. For example, who should be promoted: the marketing manager who has presided over excellent sales growth for a cutting-edge product in a hot market segment or the marketing manager who has produced modest sales growth for a product near the end of its life cycle in a shrinking market segment? To what degree is each product's sales a result of the manager's abilities and to what degree is it a consequence of the situation? Indeed, corporations regularly promote workers with strong records of achievement into jobs that require new skills only to see them fail (Fairburn & Malcomson, 2001; Kim & Aga, 2011; Lazear, 2004).

Although decisions like these are common in organizations, anecdotal evidence suggests that managers are not able to properly account for situational influences on performance when they make predictions regarding the likelihood of success of their employees in a new job (Gino & Pisano, 2011). This anecdotal evidence is supported by psychological research on attribution

which suggests that people have trouble making this type of attribution accurately (Gilbert, 1994). Attributional judgments are predictably impaired by the correspondence bias (Jones, 1990), a tendency so widespread that Ross (1977) declared it the fundamental attribution error. One way in which organizations make this error is by assuming that performance in one domain will predict performance in another domain. Kolev and Hogarth, for instance, present evidence suggesting that CEOs who are good at playing golf get paid more. In fact, however, they actually perform worse in their jobs, perhaps because they are spending too much time on the golf course (Kolev & Hogarth, 2008).

Another way in which organizations fall victim to this bias is when they rely on dispositional attributions for what are really situational effects. It leads managers to explain workers' performance in terms of their inherent abilities despite the undeniable influence of their social surroundings and institutional constraints (Mitchell & Kalb, 1981; Mitchell & Wood, 1980). This natural tendency may have contributed to interest in the search for the individual dispositions that make the best leaders (Fleishman, 1953; Ghiselli, 1971; Stogdill, 1974), a research approach that is now regarded as misguided (Davis-Blake & Pfeffer, 1989). There is, of course, a parallel literature in the field of entrepreneurship that sought to identify the traits that predispose individuals to entrepreneurial success (Hornaday & Aboud, 1971; Timmons, 1978). This too is generally regarded as having been a failure (Gartner, 1989; Shaver & L. R. Scott, 1991). Both research agendas sought to identify the distinctive personality characteristics that contribute to success and both have had disappointing results. The alternative approach assumes that success is determined in powerful ways by social settings and the opportunities they afford or the constraints they impose on leaders, entrepreneurs, and the rest of us (Burt, 1992; Gladwell, 2008; Granovetter, 1995).

The correspondence bias

How should organizations assess employee dispositional ability given noisy performance data that is confounded with situational difficulty? The answer comes from Kurt Lewin's (1931)

attributional equation: $Behavior = f(Disposition, Situation)$. In other words, an individual's performance is a joint function of both the individual's disposition (traits, abilities, and predilections) and the situation in which the individual finds him- or herself. Gilbert (1994) proposed the following simplified specification of Lewin's attributional equation: $Behavior = Disposition + Situation$. Or if our goal is to understand the individual's disposition: $Disposition = Behavior - Situation$. Gilbert's simplification holds when performance is the sum of an individual's ability and task ease. Using this equation, we can infer an individual's disposition so long as we can measure both their performance and how that performance was affected by the situation. Let's say, for example, that an Olympic figure skating judge is exactly two tenths of a point more lenient than the other judges. Any time the lenient judge is on the panel, those skaters' performance received a better evaluation. Computing the scores that skaters deserved to receive is easy: simply take two tenths of a point off from the lenient judge's rating.

The correspondence bias interferes with the simple logic of the attributional equation by ascribing too little influence to the situation and too much to the individual's disposition (Davis-Blake & Pfeffer, 1989; Jones & Harris, 1967). For instance, this bias would result in skaters being judged as better after having been evaluated by lenient judges. There is evidence that managers do indeed tend to favor dispositional attributions for employee performance (Martin & Klimoski, 1990; Mitchell & Wood, 1980). However, these studies cannot specify whether such attribution represents accuracy or bias because the research designs did not afford accurate assessments of the actual causes of performance. The laboratory evidence on the correspondence bias, for its part, is constrained by a different issue. Because no participant has objective information on both performance and situational influences on performance, it may be unrealistic to expect that anyone will be able to make the optimal judgment. Indeed, insofar as people have better information about performance than about the situation that gave rise to it, it should be no surprise that they neglect to consider the situation.

Consider, for example, the classic study by Ross, Amabile, and Steinmetz (1977). In it, participants were randomly assigned to the roles of quiz-master, quiz-taker, or observer. Quiz-masters came up with quiz questions to which they knew the answers. The quiz-takers, on the other hand, did not know the answers to all of the questions. Observers, who only observed one quiz-master and one quiz-taker, faced the difficult task of determining which of the two was more knowledgeable. On average, observers rated the randomly assigned quiz-masters as being more knowledgeable, suggesting that they based their judgments on behavior without accounting for the impact of the role. But to suggest that observers should have ignored their observations and simply rated quiz-master and quiz-taker equally knowledgeable is unrealistic because it neglects the information present in the quiz game. If the questions were all easy and the quiz-taker knew the answers to none of them, then it would not be crazy to infer that this particular quiz-master was indeed more knowledgeable than this particular quiz-taker. The fact that the quiz-master and quiz-taker must be equally knowledgeable, on average (thanks to their random assignment to roles), is of little help determining whether a particular quiz-master is more or less knowledgeable than a particular quiz-taker.

In order to assess the strength of the situational advantage conferred on quiz-masters, individual participants would have to know what percentage of questions the average quiz-taker got right in this situation. But since participants in correspondence bias experiments do not get the information that would allow them to assess the strength of the situation, perhaps it should come as no surprise that the situation gets neglected in their attributional calculus (Gilbert & Malone, 1995; Jennings, 2010). In this paper we offer a new approach in which participants have excellent information about situational influences. They are asked to make inferences about a target's disposition using information about situation and behavior that is directly comparable and in equivalent terms. This approach allows us to advance our understanding of the correspondence bias by eliminating the concerns of demand and prior beliefs while also enabling us to calculate the quantitative degree to which they neglect this information.

The present data

In this paper we test for the correspondence bias in selection decisions and ask how easy it might be to correct it. Previous research has shown that college students facing simulated admissions tasks in the lab favored applicants from favorable situations (Moore, Swift, Sharek, & Gino, 2010). We present three studies that dramatically extend the implications of prior evidence by showing that attribution errors in selection decisions impact decisions made by experts, in field settings, and across domains. In each, we examine situations in which the difficulty of a prior environment influences candidates' performance and consequently affects how they are evaluated by decision-makers. First, we examine the decisions made by admissions professionals asked to differentiate between academic excellence and lenient grading in a controlled experimental study. Next, we extend the lab paradigm to the business context and test the decisions made by executives evaluating the performance of employees in varying business climates. Finally, we return to the admissions context to present the results of an analysis of archival data documenting the decisions made by graduate business schools.

In each of these studies decision-makers are faced with the challenge of evaluating candidates' aptitude (or disposition) by observing performance and attempting to account for the varying situations. We hypothesize that, consistent with the correspondence bias, candidates from favorable situations will be evaluated more positively and selected more frequently than their counterparts with equally commendable dispositions.

In order to illustrate the analysis that tests this hypothesis, let us consider an example. Imagine that you are considering promoting one of two executives. Executive A oversees Division A, with sales of \$100 million. Executive B, on the other hand, oversees Division B, which has sales of \$110 million. Executive A brought Division A up from \$90 million while Executive B presided over Division B's increase from \$108 to \$110 million over the same period of time. It probably makes more sense to consider each executive's performance as a change from previous performance. Every dollar in sales that the divisions were already making before

the executives showed up should reduce the amount of credit they get for sales accrued during their tenure. We expect, then, that nominal performance will matter. It should, however, be discounted by the ease with which it was achieved, and this discounting should be roughly equal in size but opposite in sign. An assessment of the two executives should heavily weight the change in sales each has achieved, and should almost completely discount the total sales each wound up delivering.

Study 1 – Admissions Professionals in the lab

We put experienced admissions professionals into a fictional version of their professional roles of evaluating applicants and deciding whether to admit them. To test whether the candidate's success in the admissions process was predicted by the grading norms of their alma maters, we gave participants a simplified admissions decision task with two main pieces of information on each of nine candidates: their GPAs and the distribution from which each GPA came. Each factor was manipulated in three levels, creating a 3 (GPA above average, average, or below average) X 3 (Average GPA at alma mater: high, medium, or low) full factorial design within subjects. Average GPAs were operationalized as 3.6 (high), 3.0 (medium), or 2.4 (low) to generally correspond to the range of grading norms observed in U.S. undergraduate institutions. Performance relative to these means was operationalized as candidate GPAs that were equal to the average, 0.3 points above it, or 0.3 points below it. Because this design creates artificially round numbers, each figure in the stimuli was randomly adjusted to within ± 0.02 of the specified value.

Participants knew that the nine fictional institutions from which the nine applicants came did not differ in quality or selectivity. Grading norms at the different institutions were therefore uninformative with respect to student quality. We predicted that our expert participants would rightly reward candidates with above-average GPAs with more positive evaluations and higher probability of admission. The more interesting hypothesis is that there would also be a positive main effect of average GPA despite the irrelevance of this factor to the underlying aptitude of the

applicant. This hypothesized effect would be consistent with a judgment process that neglects to sufficiently discount GPAs in light of the ease with which they are achieved.

Method

Participants. Participants were 23 members of the staff of the undergraduate admissions office at a selective university in the United States. Participants were recruited by e-mail to a website through which the entire study was conducted. As compensation, participants were entered into a lottery in which they stood a 1 in 10 chance to win \$100. Of the participants, 61 percent were male, the average age was 34 years, and the mean professional experience in college admissions was over 5 years.

Procedure. Participants read the following instructions:

“In this exercise, you will be playing the role of a member of the admissions committee at a selective MBA program. You are selecting students who would like to obtain masters degrees in business administration. Your most important goal is to select the best candidates from among the applicants. In general, you usually have space to admit about half the applicants.

You will see the applications of nine hypothetical students. The set of applicants that you will review all graduated from colleges of similar quality and selectivity. Please review each applicant carefully in order to assess the quality of their prior academic performance in college. Please review one candidate at a time. Answer the questions about each candidate before turning the page to read about the next candidate.”

Participants received one page of information about each candidate from which to make their judgments. The information presented for each candidate including the candidate's GPA and the average GPA at their alma mater. To emphasize the candidates' performance relative to their peers, the difference between their GPA and the average was presented in the statement, “Candidate 1's GPA is [above / below] the college average by [difference].”

To add richness to the stimuli, participants also read a short transcript of ten recently completed classes. For each of the classes, the candidate's grade was listed along with the

average grade given in that class. Grades were generated for each class such that the mean of the candidate's grades were equal to their GPA and the mean of the average grades was equal to the college's average GPA. The individual class grades for each candidate were generated such that the variance was constant across candidates. The names of the fictional colleges and the names of the courses were randomized to avoid creating any confounds. Candidates were presented to each participant in a randomly determined order which was then rotated across participants in a Latin-squares design.

Participants made two judgments concerning each candidate: how successful the candidate had been in college on a 7-point scale from 1 ("very unsuccessful") to 7 ("very successful"), and how likely they would be to admit them (as a probability between 0% and 100%). After reviewing and evaluating each candidate individually, participants were asked to select four of the nine candidates for admission. Participants were permitted to review summary information for each candidate as they made their admissions decisions.

After participants finished the experimental portion of the study, they completed a small number of questions assessing their beliefs about how grades should be interpreted in the admissions process. They also told us their ages, years of experience, and job titles.

Results

While each participant produced three responses for each candidate, the results are virtually identical across measures (see Table 1). Accordingly, the three measures were standardized and averaged to create a single measure of the evaluation of each candidate (alpha reliability = .82).

This aggregate measure was then subject to a 3 (GPA) X 3 (average GPA at undergrad institution) within-subjects ANOVA. As expected, the results show a strong main effect of GPA, $F(2,44) = 85.09, p < 0.001, \eta_p^2 = 0.80$. Those candidates with GPAs above their local averages were much more likely to be admitted (70%) than those with below-average GPAs (12%).

Counter to normative expectations, but consistent with our hypothesis, the average GPA at the candidate's home institution also drove a significant main effect, and even explained as much of the variance in evaluations, $F(2, 44) = 94.24, p < 0.001, \eta_p^2 = 0.81$. This main effect suggests that it is just as advantageous to come from a school with lenient grading as it is to be above average. Indeed, the effect on the probability of admission is very similar (see Table 2). Candidates from the schools with harsh grading were admitted 12% of the time while those from the schools with lenient grading were admitted 72% of the time. Note that these differences exist despite the same distribution of performance among applicants from each "type" of school.

The interaction between GPA and Average GPA was also statistically significant, but does not explain as much variance in evaluations as the main effects, $F(4, 88) = 5.88, p < 0.001, \eta_p^2 = 0.21$. This interaction describes the fact that a better-than-average GPA had its smallest benefit for those who came from institutions with tough grading and low average grades (see Table 2).

Participants also answered two questions intended to shed light on the use of GPA statistics in real admissions decision making. We asked, "For approximately what percentage of candidates is information about their home institution's average GPA available?" While our experiment afforded decision makers access to full information about the candidates' individual performance as well as the situational influences in their environment, we suspected such data was not as readily available in the real world. Indeed, this sample of professional admissions staffers estimated that average grade information was available in only 61.5% of cases. This suggests that the magnitude of the bias measured here is likely to be conservative relative to a naturalistic process, since we provided them with more complete information than they actually have in their jobs.

In response to the question "When considering information about a candidate, how important is the average GPA at a candidate's home institution compared to their individual GPA?," the mean response (3.8) was closest to the midpoint "Equally important as the

candidate's GPA" of the scale which ranged from "Only the average GPA is important" to "Only the candidate's GPA is important." This is consistent with what we believe to be the normative decision rule, since the individual's nominal performance and the ease of the situation ought to receive equal and opposite weighting when making dispositional inferences.

Discussion

The results of these hypothetical admissions decisions made by experienced professionals suggest that candidates who happen to graduate from schools with more lenient grading norms may actually have a better chance of being accepted to college or graduate school. This is true independent of their personal performance in that situation. For example, high-performing applicants from low-GPA schools were given lower ratings than under-performing applicants from high GPA schools. This should be cause for concern among stakeholders at institutions with tougher grading norms. The high standards to which the students at these institutions are held may be mistaken for poor performance.

Despite the presence of unambiguous information on the key terms in Lewin's attributional equation, our expert participants in Study 1 were either unwilling or unable to make admissions judgments solely on the basis of applicants' dispositions. It is possible that the participants inferred variation in school quality from the variation in average grades – knowing that this relationship is true to some extent in the real world. We should point out however, that one of the few instructions for the task explicitly described the applicant's alma maters as being of similar quality. Participants may have inferred variation where they were told there was none, but this is not the most parsimonious explanation of the rather dramatic effects we observe. Instead, we argue that the results demonstrate the difficulty people – even experts – have discounting nominal performances in light of situational influences.

Study 2 – Executives and Promotion Decisions

While the decisions made by admissions professionals were consistent with our hypotheses, we fear that these effects are not isolated to university admissions. To demonstrate

that the effect is not driven by idiosyncratic properties of the admissions domain, we constructed an analogous decision task in the corporate domain. Study 2 seeks to investigate the effect of varying situational difficulty on the evaluation of employees for promotion. Participants, all of whom were working professionals with years of experience with selection and promotion decisions, were asked to evaluate twelve candidates who were up for promotion to a senior management position. We varied the difficulty of the situation and the individual performance of the candidates in order to examine how the situation affected the dispositional evaluations made by our participants. We hypothesized that managers would favor candidates benefiting from the positive light of easy business situations just as admissions had done for candidates from lenient grading environments.

Method

Participants. Participants were 129 executive education students (38% female) at a research university in the western United States participating as part of a class exercise. They had on average, 18.0 (SD=8.6) years of work experience and been involved with an average of 14.2 (SD=29.6) promotion decisions.

The experiment employed a 2 (situation: easy vs. difficult) X 3 (individual performance: low, medium and high) within-subjects design with two candidates in each cell, resulting in twelve candidates for our participants to evaluate.

Participants evaluated each of the twelve candidates, rating their performance and their suitability for promotion. Participants read a business scenario in which they are the CEO of an airline deciding whom to promote to a senior management position from a set of twelve candidates. The most important measure of their performance is the percentage of flights that leave on time from each of their airports. Moreover, of the twelve airports managed by the candidates, there were two types of airports: Half were historically quite punctual and half that ran late more often.

Task difficulty was manipulated via the historical on-time percentage of flights at each airport. Hard airports had 70% of flights on time; easy airports had 85% on-time flights. Individual performance was manipulated relative to this situational effect: low performance was 5% below the average, medium performance was at the average and high performance was 5% above the average. We varied these artificially round numbers by a few tenths of a percentage point to increase the realism of the scenario.

In order to ensure that candidates were being evaluated solely on the basis of their on-time performance and not by airport, we counter-balanced each of the possible match-ups between airport and performance and randomized the order in which participants encountered the candidates.

For each candidate, participants were shown the airport's historical on-time performance across ten years: five years preceding the candidate's tenure and the most recent five years under their management. The five-year average on-time percentage was presented with each block of data. Participants answered two questions about each candidate: (a) Performance: "*Please rate the performance of Candidate Name, the manager in charge of operations at this airport.*" on a 7-point scale anchored at 1 (*Very Bad*) and 7 (*Very Good*), and (b) Promotion Worthiness: "*Is Candidate Name worthy of promotion?*" on a 7-point scale anchored at 1 (*Definitely No*) and 7 (*Definitely Yes*).

Results

Table 3 summarizes the descriptive statistics for the dependent measures, *Performance* and *Promotion Worthiness*. We averaged the ratings given to the two candidates in each cell for the two measures and submitted them each to a 2 (difficulty) X 3 (performance) repeated-measures ANOVA.

The results for the performance rating revealed a significant within-subjects effect of both situational difficulty, $F(1,128) = 80.40, p < .001, \eta_p^2 = .39$ and actual performance, $F(2, 256) = 301.92, p < .001, \eta_p^2 = .70$ on the candidates' performance ratings. The interaction between our

two main effects, difficulty X performance, was not significant, $F(2, 256) = 0.692, p > .05$, indicating that while our participants gave higher ratings to candidates from easy airports, they were consistent in their ratings within each difficulty category.

The results for the promotion worthiness rating followed a similar pattern, with significant within-subjects effects for both difficulty, $F(1, 128) = 73.71, p < .001, \eta_p^2 = .37$, and performance, $F(2, 256) = 248.539, p < .001, \eta_p^2 = .66$, and a non-significant interaction of difficulty X performance, $F(2, 256) = 0.480, p > .05$.

Discussion

If our participants had not been affected by the correspondence bias, then candidates working at easier airports should not have been awarded higher ratings than their dispositional peers working at more difficult airports. However, what we find is that those who are fortunate enough to be assigned to easy airports are evaluated significantly more positively. We can make strong conclusions with respect to the normative status of participants' decisions, thanks to the control of the experimental paradigm. The experienced professionals who served as participants in the first two studies made consistent errors in deciding whom to admit or promote. They favored those from easy situations, rewarding the fortunate and punishing the unfortunate.

However, legitimate questions remain regarding the degree to which they actually make these same errors in their work. One might be tempted to predict that the effects we observe in the first two studies would be even more dramatic in professional domains that do not afford such clear information about the size of situational effects. On the other hand, it may be that organizations can provide some correctives for individual biases (Heath, Larrick, & Klayman, 1998). Hypothetical decisions in an experiment could be limited in their validity as a model of fully consequential and incentivized decision making. We seek to address these limitations by demonstrating convergent evidence in a sample of analogous decisions made in the field.

Study 3 – Actual Admissions Decisions

Studies 1 and 2 demonstrate the difficulty, at least in the lab scenarios we created, of properly incorporating information about the difficulty of the situation when evaluating performance. Of course, even with a sample of professionals one can always ask whether the motivation and consequences of real decisions can be emulated in a laboratory design. Moreover, it is possible that the effect we observe in carefully controlled experimental settings shrinks to insignificance relative to all the other influences in the complexity of real admissions decisions. To test our hypotheses in the full light of day, we move to an archival sample of actual admissions decisions.

Method

We obtained admissions data from four selective (admission rates between 16% and 39%) MBA programs in the United States. These data include each applicant's undergraduate GPA and alma mater, as well as a number of useful control variables such as age, gender, race, national citizenship, years of work experience, standardized test scores (in this case, the GMAT test), and interview performance. Each school provided multi-year data ranging from three to eleven years of coverage and drawing from years between 1999 and 2009. Annual applicant totals per school varied from 1,236 to 3,742, yielding between 5,700 and 20,100 applicants per school across all available years, and a total dataset of 56,796 applicants. Of this set, 24,994 were graduates of foreign undergraduate institutions, and are excluded from the analyses due to incompatible grading systems.

We collected additional data describing the undergraduate institutions (i.e. school quality, average GPA) from which the applicants had graduated. We began with two measures of school quality. The first, Jack Gourman's ratings of undergraduate institutions is an academically cited continuous measure of the quality of education provided for nearly every U.S. institution (Gourman, 1989). As an additional measure of quality, we used the US News & World Report annual ranking of the US colleges and universities (McGrath, 2004). While the US News measure of quality is not academically reviewed and does not offer complete transparency in its

methodology, it captures some dimensions of school quality – such as perceived prestige and alumni success – that may be important in our model. It is possible that perceptions of institution quality are actually more direct drivers of admissions decisions than the underlying quality Gourman strives to capture. As a final, more objective measure of school quality we used the average GMAT score of entering students.

We attempted to collect data on average grades for all institutions represented in our sample of applicants. The data on average grades we collected came from a variety of sources. The largest single source was a publicly available database of average GPAs assembled for the purpose of documenting trends in grade inflation (Rojstaczer & Healy, 2010), but this source was not comprehensive. The remaining institutions for which we lacked average grade information were contacted directly by email, or when that failed to produce a response, by telephone. Many declined to provide us with information about average grades. In the end, we were able to obtain average GPA data from 198 institutions, corresponding to 20,913 applicants or 77% of the domestic applicants in our sample. Average GPA of the applicants' undergraduate institution was the most restrictive variable in our dataset, with the other predictors ranging from 85% - 100% complete.

Applicants' performance was measured by their deviation from the average GPA at their undergraduate institution, henceforth relative GPA. By modeling performance as the difference between those two raw data sources, we were able to directly test our hypothesis that applicants' situations affect their outcomes above and beyond their unique performances in that situation. Relative GPA is a standard unit of performance that can be used to compare the disposition of students across situations. We are aware of the problems that can accompany the use of difference scores, but we do not believe our analyses are threatened by the problems documented in the literature (Edwards, 2001; Peter, Churchill, & Brown, 1993). While a difference score composed of two scale measures will exacerbate reliability problems, the difference taken in our analysis is of two known quantities for which reliability is a smaller concern for our purposes.

Also, including both GPA and average GPA in a model complicates their interpretation since average GPA (together with the student's performance) is a component of one's nominal GPA. Using relative GPA also avoids problems with logistic regression assumptions because the distribution of nominal GPAs is truncated at 4.0.

Interview ratings were standardized within each application year at each admitting school to account for different systems of evaluation. Average GPA at a candidate's undergraduate institution was mean centered before being included in the model to aid in interpretation. Racial identifications are presented as comparisons to the baseline category of Caucasian which was the largest group in the sample.

Results

The means, standard deviations, and bivariate correlations between applicant level variables are shown in Table 4, while Table 5 reports the same statistics for school level variables. As shown in Table 4, there is, not surprisingly, a positive relationship between applicants' undergraduate grades and admission decisions: Those who were admitted had higher GPAs ($M = 3.35$) than those who were not ($M = 3.29$), $t(36384) = 13.12$, $p < .001$. It is possible that the positive effect of institutional GPAs is because better universities have higher average grades. Indeed, average grades tend to be positively related to institution quality with correlations between average GPA and other measures of quality ranging from $r = .18$ for the Gourman ratings to $r = .62$ for the US News ratings. Average absolute GPA may in fact be correlated with student ability. A more complete picture, then, is provided by a full multiple regression.

In order to examine which factors predict admission, we conducted a series of ordinal logistic regressions, modeling the outcomes of being waitlisted and accepted as ordered improvements over the reference group of rejected applicants. Table 6 summarizes these models, showing how the average GPA and applicants' relative GPA influence admissions outcomes. In each of these models, we hypothesize that applicants' relative GPAs will be

positively associated with admission indicating that above average performances are rewarded, but that the average GPA at the applicant's alma mater will also be positive and significantly related to admission. This would represent an error, since, holding all else equal, graduate schools should not want to reward candidates solely for graduating from institutions with more lenient grading norms.

Models 1 through 3 build an increasingly complex model of admissions decisions starting with only our central variables of interest, and then adding the relevant academic and demographic terms. In all three models, relative GPA is positive and significant, indicating that individual performance is rewarded by increased rates of acceptance. Average GPA is significant and positive as hypothesized in model 1, but this simple model explains a very small amount of the variance in admissions outcomes. In model 3, the Average GPA term is not significant. This is likely due to the fact that models 1 through 3 do not control for the differences in application year and admitting school. Admission rates vary dramatically across these situations, from 16% to 39% by school, and from 20% to 33% by application year. Subsequent models address this shortcoming.

Models 4 and 5 introduce dummy variables for the four admitting schools and eleven years in which admission decisions were made, while model 6 uses both factors to control for the significant variation in baseline acceptance rates. Model 6 shows that both a candidate's relative GPA and the average GPA at their alma mater have a significant positive effect on their likelihood of acceptance.

Model 7 replicates the findings of model 6 but drops the candidate's interview scores from the equation. After including all of the control variables available in our dataset, the interview score restricts our sample size more than any other control. Dropping it in model 7 allows us to double our sample size and maintain a representative sample through more restrictive analyses. Comparing model 7 to model 6 reveals no qualitative changes in any of the

predictors although the relative influence of average GPA increases compared to relative GPA. Further models discussed will generally be adaptations of model 7.

Table 7 presents models 8 through 11, which demonstrate the robustness of our results to three different measures of applicant's undergraduate school quality: US News ratings, Gourman (1989) ratings, and the average entrance SAT scores of matriculating students. In order to compare the relative influence of these three predictors, each was standardized. The relationship between school quality and applicant's admission is critical to the interpretation of our results because of the correlation between quality and average GPA. It is tempting to consider whether the large positive effect of average GPA is related to the fact that many of those students benefiting are from good schools with high grading norms. Models 8-11 show that no matter how school quality is measured, average GPA has a distinct and positive impact on candidate outcomes.

In each of the other models presented, only the US News ratings are included as the measure of school quality for two reasons. First, of the three measures, US News ratings are most closely related to the public reputation of the schools. It is this reputation that we expect influences the evaluation of their alumni more than the less visible educational values captured by the other measures. The data support these inferences. While each of the quality predictors are significant when included alone, model 11 shows that effect of school quality as measured by the US News ratings are more stable in both significance and magnitude.

Table 8 presents three new models to address alternative analysis approaches and concerns. Model 7 is presented again for ease of comparison. Model 12 includes an interaction between the two primary predictors, relative and average GPA. The interaction is significant and positive, but it does not affect the results from the main effect which remain positive, significant, and of similar magnitude. We do not emphasize this finding due to the complicated and sensitive nature of interactions between two continuous variables in logistic models.

Model 13 tests the analysis with a subset of the sample of applicants. While we maintain that it is normatively ideal for admissions decision makers to discount applicants' nominal GPAs to account for variation in grading norms, this may be difficult to do in practice. It was quite difficult for us to get access to average GPA data from the wide selection of undergraduate institutions represented in our sample of applicants. Our sample of professional admissions workers from Study 1 indicated that they had some information about grading norms for nearly two thirds of applicants. Since it is likely that the top undergraduate institutions are better known to admissions staffers, we restricted the sample in model 13 to those applicants who graduated from schools in the top quartile of US News ratings. This decreased the sample by only 35% compared to the model 7 because so many applicants in the sample graduated from these top institutions. The model shows that the applicant's GPA continues to be significant and positive, but that even in this top quartile, average GPA is robustly significant and positive. In fact, in this subset, the magnitude of the effect of average GPA is slightly larger than that of the candidate's relative performance. All else equal, applicants who are a tenth of a point above average at their school are no more likely to be accepted than applicants that had average GPAs at a school where the grading norms were a tenth of a point higher than other schools.

Finally, in model 15 we ensure that our conclusions are robust to an analysis without our computed relative GPA measure. In this model, normative admissions decisions would discount situational influences by weighting average GPA negatively and to a similar degree that nominal GPA is positive. We see that average GPA is not significantly different than zero, consistent with the conclusion that decision makers are not adjusting nominal performance in light of information about the situation.

Discussion

We argue that the tendency for admissions decisions to ignore grading leniency is explainable by attributional processes that take performance at face value and fail to discount it by situational influences. The result is that applicants from schools with lenient grading are

evaluated more positively simply because their GPAs are higher. Model 8 shows that applicants benefitting from an average GPA one standard deviation above the mean (just 0.17 points) are 31% more likely to move up from denied to waitlisted or from waitlisted to accepted. Figure 1 illustrates the magnitude of this effect.

However, the field data do not allow us to rule out two viable alternative explanations for our results. First, it might not be fair to assume that admissions staffs have perfect information about average grades at all other academic institutions. After all, these data were costly for us to collect. Without the data, it would be difficult for admissions departments to use the information to discount grades appropriately, and they would simply have to rely on absolute GPAs as useful (if imperfect) measures of academic performance. On the other hand, this information is critical to their ability to interpret applicants' GPAs, and admissions offices are likely to have better ways of obtaining these data than did we. A failure to demand it, in and of itself, suggests a failure to appreciate its value.

Second, it is possible that lenient grading captures something important about the undergraduate institution. While we considered the roles of institution quality, public/private status, tuition, and student quality as control variables, it is possible that lenient grading covaries with other important features of institutions that are not captured in our control variables. The experimental data from the first study can help us to address this concern to some degree. Professional admissions workers evaluated candidates from various fictional schools whose grading leniency varied substantially but that were equal in quality and selectivity. This makes it unlikely that our participants favored candidates from lenient-grading institutions because they believed that the students were actually better at lenient-grading institutions.

General Discussion

Many studies in the social psychology and organizational behavior literatures have found that people tend to attribute too much influence to disposition and too little influence to situational factors impinging on the actor when explaining others' behaviors, achievements, and

failures. This common tendency, labeled the correspondence bias or the fundamental attribution error, has been shown to be robust across a variety of contexts and situations. Yet most of the evidence about this bias comes from laboratory experiments with college students as participants, and its implications for field settings and organizational outcomes are seldom examined. Using data from both the experimental laboratory and the field, we extend prior research by investigating whether this tendency leads experienced professionals to make systematic mistakes in their selection decisions, favoring alumni from academic institutions with lenient grading and employees working in favorable business climates. Our results indicate that candidates who have demonstrated high performance thanks to favorable situations are more likely to be rated highly and selected. Across all three studies, the results suggest that experts take high performance as evidence of high ability and do not sufficiently discount it by the ease with which that performance was achieved. High grades are easier to achieve in a lenient-grading environment and so are less indicative of high performance than are the same grades that were earned from an institution with tougher grading standards. Sky-high on-time percentages should be less impressive at an airport that was running well before the manager got there. While we focused on two selection scenarios, we believe the results speak to other selection and evaluation problems.

Indeed, we see consistent evidence of situation neglect in contexts where political and business leaders are credited with performance that derives directly from stochastic economic factors. Voters face a Lewinian dilemma when they evaluate the performance of incumbent politicians running for re-election. They should reward politicians who create positive change for their constituencies while considering what portion of those changes were due to lucky or exogenous factors. Wolfers(2002) finds that voters, like our admissions professionals and executives, favor politicians that had the good luck to work under favorable conditions. Voters are more likely to reelect incumbents after terms marked by positive national economic trends or (in the case of oil-rich states) high oil prices. CEOs also benefit from fortuitous economic

conditions for which they are not responsible. Bertrand and Mullainathan (2001) present evidence that CEO compensation is driven to equal degrees by their management and the uncontrollable economic conditions in which they managed. Stakeholders in these cases have strong incentives to reward leaders who add value above the vagaries of the economy, but they seem blind to the difference.

It is often the case that structural and situational factors are the most powerful influences on behavior. Within organizations, for example, it is easier to succeed in some jobs than in others (Pfeffer & Davis-Blake, 1986). Sometimes people will achieve positive outcomes simply because of a beneficent environment. It is easier to achieve success as a manager when your team is strong than when your team is weak. Likewise, it is easier to obtain a strong education in an excellent private school than in an under-funded public school. And it is easier to achieve high grades at schools where everyone is graded leniently. So it would be a mistake to neglect situational effects on performance, but that is what our data suggest that even experts and professionals tend to do.

Organizations often adopt practices because they are legitimate, popular, or easy to justify (W. R. Scott, 1998; Staw & Epstein, 2000). That may help explain why we observed such consistency in admissions policies in neglecting to consider differences in grading leniency between institutions. On the other hand, it is reasonable to assume that finding a better system to evaluate applicants would improve admissions decisions, allowing the schools that do it to identify strong candidates that other schools neglect. The Oakland Athletics baseball team did just this when it pioneered a new statistical approach to identifying promising baseball players to recruit (Lewis, 2004). Their success has since been emulated by other teams, fundamentally changing the way baseball's talent scouts pick players. However, the problem for admissions departments may be somewhat more complicated because explicitly tarring some institutions as lenient-grading is likely to elicit energetic protests if they ever find out about it (e.g. Woo, 1997).

It is common in organizations for the abilities of an individual, a department, or a division to be shrouded in complicating or confounding influences that make them difficult to detect or measure (Siggelkow & Rivkin, 2009). Indeed, as much as ratings systems like grades and performance metrics like on-time percentages can help clarify standards for evaluation, they can also be used to obscure performance (Fleischer, 2009). Variation in grading standards between institutions obscures the value of using grades to measure student performance. It is probably in the interest of lenient-grading institutions to hide the degree of their leniency. Consistent with this motive, recent years have seen changes in the disclosure that institutions are willing to make (Marklein, 2003). Fewer academic institutions are willing to disclose average grading data or class rankings for their students or alumni. When we contacted institutions to inquire regarding average grades elite, expensive, private institutions – those with the highest average grades – were most likely to decline to disclose the information.

Organizational Image, Legitimacy, and Stakeholder Appraisals

The strategic use of scoring and assessment metrics has implications at the organization level because of the way that institutions compete. Scott and Lane advanced a theory of organizational image in which stakeholders (both members as well as outside audiences) play a key role in shaping the organization's image by making legitimacy appraisals that can counterbalance the organization's attempts at image management (2000). This model is built on the dual premises that organizations and their members derive personal and economic benefits from promoting a positive image (Ashforth & Mael, 1989; Fombrun & Shanley, 1990), but that salient audiences have a role in validating that image (Elsbach & Sutton, 1992; Schlenker, 1985). These forces form an equilibrium that balances the organization's incentives for an unbounded positive spin with the utility gained by stakeholders from an image grounded in reality. Scott and Lane (2000) term the specific mechanism by which this equilibrium is reached *reflected stakeholder appraisals*. In the present paper we have investigated a setting in which

stakeholders may have difficulty judging the appropriateness of image-relevant information which could then threaten the stability of the reflected stakeholder appraisal equilibrium.

In the context of higher education, graduating students are among the primary interfaces through which employers, graduate schools, and communities interact with undergraduate institutions. Their reputation in the form of grades contributes to the reputation (Gotsi & Wilson, 2001) of the organization. As such, undergraduate institutions have an incentive to promote an image of intelligence and achievement to these outside audiences by maintaining a relatively high grade distribution. Given the tremendous value of being able to place alumni in better graduate schools and in better jobs, universities cannot be expected to go too far in seeking to curtail grade inflation. For example, universities are unlikely to implement meaningful institutional changes such as replacing grades with percentile rankings. Instead, we should expect academic institutions to pay lip service to the importance of high academic standards while at the same time avoiding publicizing average grade distributions and avoiding reporting class rank data on their students.

Do we see unchecked escalation of grade distributions by a market full of organizations unconstrained by the critical feedback from shareholders? Of course there are multiple mechanisms supporting a moderate equilibrium even without functioning shareholder criticism of the type we have described, but some data suggest grade inflation is a prolonged and significant trend in U.S. Education (Rojstaczer & Healy, 2010). More troubling perhaps are recent anecdotal reports of institutions manipulating their grade distribution with the publicly expressed intent of influencing the selection decisions of hiring firms (Rampell, 2010). Clearly, these institutions are anticipating that employers will not sufficiently discount the grades of their alumni to eliminate the advantage their inflated grades will confer.

Limitations and Directions for Future Research

Our studies are subject to several important limitations. First, the sample used in our first study was relatively small due to the size of the admissions department that participated, even

though the results were highly significant. In addition, the first and second studies employed hypothetical decisions, which may have limited validity as a model of fully consequential and incentivized decision making. Future research could benefit from a more qualitative research approach to investigate how admissions and promotion decisions are made by various organizations. As for Study 3, there are many variables (such as variations in average GPA by discipline within a school) for which we did lack information and thus could not control in our analyses. These variables may have important influences on admission decisions that are not captured in the present research. Although these are important limitations, it is also worth noting that the limitations differ across studies and yet the findings are robust.

The conclusions implied by our results as well as the limitations of our research bring forth some fruitful and interesting possible avenues for future research. One interesting question is whether other academic selection contexts would show the same patterns as business school admissions decisions. Law schools, for instance, use the Law School Admissions Council, an organization that (among other things) processes applications for law schools and provides a service that gives schools a sense of where a given applicant's GPA falls relative to other applicants that the LSAC has seen from that same institution. The Graduate Management Admissions Council does not process business school applications and so does not provide an equivalent service for business schools. Does the LSAC's assistance help law schools make better admissions decisions?

Similarly, future research could explore the implications of the correspondence bias for promotions of business professionals. Just as educational institutions vary with respect to the ease of achieving high grades, so do companies, industries, and time periods differ with respect to the ease of achieving profitability. There are some industries (such as airlines) that are perennially plagued by losses and whose firms have trouble maintaining profitability. There are other industries (such as pharmaceuticals) that have seen more stable profitability over time.

And clearly there are changes over time in industry conditions that drive profitability; for example, global oil prices drive profitability among oil companies.

We believe an important avenue for further investigation lies in continuing the study of the correspondence bias in empirical settings with organizationally-relevant outcomes. A more thorough understanding of the implications of this common bias for organizations could be achieved by further investigating business decisions such as promotions. There are also a multitude of other business decisions in which a latent variable of interest is seen in the context of varying situational pressures. Investment returns, sports achievements, and political success are all domains in which judgments are vulnerable to the tendency to insufficiently discount the influence of the situation. We expect that the correspondence bias affects outcomes in these domains.

Our theory holds that a firm's good fortune (in the form of greater profits) will be mistaken as evidence for the abilities of its managers. If this is so, then we should more often see employees of lucky firms being promoted than of unlucky firms (Rosenzweig, 2007). We would expect, for instance, that pharmaceutical executives are more likely to be hired away to head other firms than are airline executives. However, this finding might be vulnerable to the critique that pharmaceutical executives actually are more capable than are airline executives—after all, their firms are more consistently profitable. Therefore, a better way to test this prediction would be using an industry (such as oil) in which fortunes fluctuate over time due to circumstances outside the control of any firm's managers. Our prediction, then, would be that oil executives are more likely to be hired away to head other firms when the oil industry is lucky (i.e., oil prices are high) than when the industry is unlucky (i.e., oil prices are low).

Theoretical Contributions

Our results contribute to the literature on the psychological process at work in comparative judgment, a literature that stretches across psychology (Windschitl, Kruger, & Simms, 2003), economics (Camerer & Lovallo, 1999), and organizational behavior (Moore &

Cain, 2007). In this paper, we extend previous research by examining judgmental contexts in which expert decision-makers are comparing outcomes that vary with respect to both nominal performances and their ease. We should also point out that these results are, in a number of ways, more dramatic than the results of previous research showing biases in comparative judgment. Previous results have been strongest when participants themselves are the focus of judgment (Moore & Kim, 2003; Windschitl et al., 2003). Biases in comparative judgment shrink when people are comparing others, and shrink still further when they have excellent information about performance by those they are comparing (Moore, 2007). Biases disappear when comparisons are made on a forced ranking scale (Klar & Giladi, 1997). In this paper, we have shown comparative judgments to be powerfully biased even when people are evaluating others about whom they have complete information (as clearly modeled in Study 1), and even when the assessments (e.g., admission decisions) are made on a forced distribution that prevent them from rating everyone as better than everyone else.

While attribution biases have been extensively studied, the vast majority of this research has been conducted with convenience samples of students. This raises the familiar concern that deviations from normative decision outcomes are the result of insufficient motivation, or that they could be corrected with sufficient experience. Indeed, some studies have found that experts often employ superior decision strategies and enjoy more positive outcomes (see Shanteau, 1988). Our results however, suggest that experts do not discount nominal performance in light of information about the situation any differently than students in a hypothetical admissions task (Moore, Swift, Sharek, & Gino, 2010). The robustness of this effect in expert populations connects a well-known psychological phenomenon to a world of selection decisions with tangible and troubling implications.

Finally, our work contributes to research on human resource and employment decisions. Previous research has examined a number of factors that influence employment decisions such as hiring and promotion. For instance, research suggests that objective applicant qualifications,

such as education or work experience, play a large role in hiring and selection decisions (e.g. Cable & Judge, 1997; Gilmore & Ferris, 1989) but are often subject to situational factors. One of the qualifications previous studies have examined is applicants' general mental ability, which is often measured using the applicants' GPA. Consistently, research has demonstrated that an applicant's GPA influences evaluation and subsequent hiring decisions (Dunn, Mount, Barrick, & Ones, 1995; J. E. Hunter & R. F. Hunter, 1984). Our results contribute to this literature by suggesting that future studies may benefit from also including information about the average GPA of the institution where the applicant studied, or other appropriate measures of situational influence.

Conclusion

Each of our studies supports the hypothesis that people rely heavily on nominal performance (such as GPA) as an indicator of success while failing to sufficiently take into account information about the distributions of performances from which it came. To the extent that admissions officers and hiring managers generally show the same biases we found, graduate programs and businesses are collectively choosing to select candidates who demonstrated their merit in favorable situations rather than selecting the best candidates. The consequences could be substantial for both the sufficiently qualified but unselected candidates as well as for the organizations that systematically select lower performing candidates than they could.

References

- Ashforth, B. E., & Mael, F. (1989). Social identity theory and the organization. *Academy of Management Review*, *14*(1), 20-39.
- Attewell, P. (2001). The Winner-Take-All High School: Organizational Adaptations to Educational Stratification. *Sociology of Education*, *74*(4), 267 - 295.
- Bagues, M., Labini, M. S., & Zinovyeva, N. (2008). Differential Grading Standards and University Funding: Evidence from Italy. *CESifo Economic Studies*, *54*, 149-176.
- Bertrand, M., & Mullainathan, S. (2001). Are CEOs Rewarded for Luck? The Ones without Principals Are. *Quarterly Journal of Economics*, *116*(3), 901–932. MIT Press.
- Burt, R. S. (1992). *Structural holes: the social structure of competition* (p. 313). Cambridge, MA: Harvard University Press.
- Cable, D. M., & Judge, T. a. (1997). Interviewers' perceptions of persons-organization fit and organizational selection decisions. *The Journal of applied psychology*, *82*(4), 546-61.
- Camerer, C. F., & Lovallo, D. (1999). Overconfidence and excess entry: An experimental approach. *The American Economic Review*, *89*(1), 306-318.
- Davis-Blake, A., & Pfeffer, J. (1989). Just a Mirage: The Search for Dispositional Effects in Organizational Research. *Academy of Management Review*, *14*(3), 385-400.
- Dunn, W. S., Mount, M. K., Barrick, M. R., & Ones, D. S. (1995). Relative importance of personality and general mental ability in managers' judgments of applicant qualifications. *Journal of Applied Psychology*, *80*(4), 500-509. doi:10.1037/0021-9010.80.4.500
- Edwards, J. R. (2001). Ten Difference Score Myths. *Organizational Research Methods*, *4*(3), 265-287. doi:10.1177/109442810143005
- Elsbach, K. D., & Sutton, R. I. (1992). Acquiring organizational legitimacy through illegitimate actions: A marriage of institutional and impression management theories. *Academy of Management Journal*, *35*(4), 699-738.

- Fairburn, J. A., & Malcomson, J. M. (2001). Performance, Promotion, and the Peter Principle. *Review of Economic Studies*, 68(1), 45-66. doi:10.1111/1467-937X.00159
- Fleischer, A. (2009). Ambiguity and the Equity of Rating Systems: United States Brokerage Firms, 1995–2000. *Administrative Science Quarterly*, 54(4), 555-574. doi:10.2189/asqu.2009.54.4.555
- Fleishman, E. A. (1953). The measurement of leadership attitudes in industry. *Journal of Applied Psychology*, 37(3), 153-158.
- Fombrun, C., & Shanley, M. (1990). What's in a name? Reputation building and corporate strategy. *Academy of Management Journal*, 33(2), 233-258.
- Gartner, W. B. (1989). Who is an entrepreneur? Is the wrong question. *Entrepreneurship Theory and Practice*, 12(4), 11–32. New York.
- Ghiselli, E. E. (1971). *Explorations in managerial talent*. Pacific Palisades, CA: Goodyear Pub. Co.
- Gilbert, D. T. (1994). Attribution and interpersonal perception. In A. Tesser (Ed.), *Advanced Social Psychology* (pp. 99-147). New York, {NY}: McGraw-Hill.
- Gilbert, D. T., & Malone, P. S. (1995). The Correspondence Bias. *Psychological Bulletin*, 117(1), 21-38.
- Gilmore, D. C., & Ferris, G. R. (1989). The Effects of Applicant Impression Management Tactics on Interviewer Judgments. *Journal of Management*, 15(4), 557-564. doi:10.1177/014920638901500405
- Gino, F., & Pisano, G. P. (2011). Why Leaders Don't Learn from Success. *Harvard Business Review*, 89(April), 68–74.
- Gladwell, M. (2008). *Outliers: The Story of Success* (p. 309). New York: Little, Brown and Company.
- Goldman, R. D., & Widawski, M. H. (1976). A Within-Subjects Technique for Comparing College Grading Standards: Implications in the Validity of the Evaluation of College

Achievement. *Educational and Psychological Measurement*, 36(2), 381-390.

doi:10.1177/001316447603600217

- Gotsi, M., & Wilson, A. M. (2001). Corporate reputation: seeking a definition. *Corporate Communications: An International Journal*, 6(1), 24-30. doi:10.1108/13563280110381189
- Gourman, J. (1989). *The Gourman report: a rating of undergraduate programs in American & international universities* (7th ed.). Los Angeles, CA: National Education Standards.
- Granovetter, M. S. (1995). *Getting a job: a study of contacts and careers* (p. 251). Chicago: University of Chicago Press.
- Heath, C., Larrick, R. P., & Klayman, J. (1998). Cognitive Repairs: How Organizational Practices Can Compensate for Individual Shortcomings. In B. M. Staw & L. L. Cummings (Eds.), *Research in Organizational Behavior* (Vol. 20, pp. 1-37). Greenwich, CT: Jai Press, Inc.
- Hornaday, J. A., & Aboud, J. (1971). Characteristics of successful entrepreneurs. *Personnel Psychology*, 24, 141-153.
- Hunter, J. E., & Hunter, R. F. (1984). Validity and utility of alternative predictors of job performance. *Psychological Bulletin*, 96(1), 72-98. doi:10.1037//0033-2909.96.1.72
- Jennings, K. E. (2010). Determining the Internal Consistency of Attitude Attributions. In R. Catrambone & S. Ohlsson (Eds.), *Proceedings of the 32nd Annual Conference of the Cognitive Science Society* (pp. 978-984). Austin, Texas: Cognitive Science Society.
- Jones, E. E. (1990). *Interpersonal perception*. New York: W H Freeman/Times Books/ Henry Holt & Co.
- Jones, E. E., & Harris, V. A. (1967). The attribution of attitudes. *Journal of Experimental Social Psychology*, 3(1), 1-24.
- Kim, P. H., & Aga, A. (2011). Leaps of faith: The nature of opportunistic generalizations across domains.

- Klar, Y., & Giladi, E. E. (1997). No one in my group can be below the group's average: A robust positivity bias in favor of anonymous peers. *Journal of Personality and Social Psychology*, 73(5), 885-901. doi:10.1037/0022-3514.73.5.885
- Koepfel, D. (2006, December 31). Those Low Grades in College May Haunt Your Job Search. *New York Times*. New York.
- Kolev, G. I., & Hogarth, R. (2008). Illusory correlation in the remuneration of chief executive officers: It pays to play golf, and well. *Working Papers (Universitat Pompeu Fabra. Departamento de Econom'ia y Empresa)*.
- Lazear, E. P. (2004). The Peter Principle: A Theory of Decline. *Journal of Political Economy*, 112(S1), S141-S163. doi:10.1086/379943
- Lewin, K. (1931). The conflict between Aristotelian and Galileian modes of thought in contemporary psychology. *Journal of General Psychology*, 5, 141-177.
- Lewis, M. M. (2004). *Moneyball: the art of winning an unfair game* (p. 317). W.W. Norton.
- Marklein, M. B. (2003). More public high schools dropping class ranking. *USA Today*.
- Martin, S. L., & Klimoski, R. J. (1990). Use of verbal protocols to trace cognitions associated with self- and supervisor evaluations of performance. *Organizational Behavior and Human Decision Processes*, 46(1), 135-154. doi:10.1016/0749-5978(90)90026-6
- McGrath, A. (Ed.). (2004). *U.S. News Ultimate College Guide 2005* (p. 1763). Sourcebooks.
- Mitchell, T. R., & Kalb, L. S. (1981). Effects of outcome knowledge and outcome valence on supervisors' evaluations. *Journal of Applied Psychology*, 66(5), 604-612. doi:10.1037/0021-9010.66.5.604
- Mitchell, T. R., & Wood, R. E. (1980). Supervisor's responses to subordinate poor performance: A test of an attributional model. *Organizational Behavior and Human Performance*, 25(1), 123-138. doi:10.1016/0030-5073(80)90029-X

- Moore, D. A. (2007). Not so above average after all: When people believe they are worse than average and its implications for theories of bias in social comparison ☆. *Organizational Behavior and Human Decision Processes*, 102(1), 42-58. doi:10.1016/j.obhdp.2006.09.005
- Moore, D. A., & Cain, D. M. (2007). Overconfidence and underconfidence: When and why people underestimate (and overestimate) the competition. *Organizational Behavior and Human Decision Processes*, 103(2), 197–213. Elsevier.
- Moore, D. A., & Kim, T. G. (2003). Myopic Social Prediction and the Solo Comparison Effect. *Journal of Personality and Social Psychology*, 85(6), 1121-1135.
- Moore, D. A., Swift, S. A., Sharek, Z. S., & Gino, F. (2010). Correspondence Bias in Performance Evaluation: Why Grade Inflation Works. *Personality and Social Psychology Bulletin*, 36(6), 843-852. doi:10.1177/0146167210371316
- Peter, J., Churchill, G. A. J., & Brown, T. J. (1993). Caution in the use of difference scores in consumer research. *Journal of Consumer Research*, 19(March), 655-662.
- Pfeffer, J., & Davis-Blake, A. (1986). Administrative Succession and Organizational Performance: How Administrator Experience Mediates the Succession Effect. *Academy of Management Journal*, 29(1), 72-83. Academy of Management.
- Rampell, C. (2010). In Law Schools, Grades Go Up, Just Like That. *New York Times*, A1. New York.
- Rojstaczer, S., & Healy, C. (2010). Grading in American Colleges and Universities. *Teachers College Record*, (March).
- Ross, L. D. (1977). The intuitive psychologist and his shortcomings: Distortions in the attribution process. In L. Berkowitz (Ed.), *Advances in Experimental Social Psychology* (Vol. 10, pp. 173-220). New York: Academic. doi:10.1016/S0065-2601(08)60357-3
- Ross, L. D., Amabile, T. M., & Steinmetz, J. L. (1977). Social roles, social control, and biases in social-perception processes. *Journal of Personality and Social Psychology*, 35(7), 485-494.

- Schlenker, B. R. (1985). Identity and self-identification. In B. R. Schlenker (Ed.), *The self and social life* (pp. 65-99). New York: McGraw-Hill.
- Schmidt, F. L., & Hunter, J. E. (1998). The Validity and Utility of Selection Methods in Personnel Psychology : Practical and Theoretical Implications of 85 Years of Research Findings. *Psychological Bulletin*, *124*(2), 262-274.
- Scott, S. G., & Lane, V. R. (2000). A Stakeholder Approach to Organizational Identity. *The Academy of Management Review*, *25*(1), 43 - 62.
- Scott, W. R. (1998). *Organizations: Rational, Natural, and Open Systems* (4th ed.). Upper Saddle River, NJ: Prentice Hall.
- Shanteau, J. (1988). Psychological characteristics and strategies of expert decision makers. *Acta Psychologica*, *68*(1), 203-215. doi:10.1016/0001-6918(88)90056-X
- Shaver, K. G., & Scott, L. R. (1991). Person, process, choice: The psychology of new venture creation. *Entrepreneurship Theory and Practice*, *16*(2), 23-45.
- Siggelkow, N., & Rivkin, J. W. (2009). Hiding the evidence of valid theories: how coupled search processes obscure performance differences among organizations. *Administrative Science Quarterly*, *54*(4), 602. SAGE Publications.
- Staw, B. M., & Epstein, L. D. (2000). What Bandwagons Bring: Effects of Popular Management Techniques on Corporate Performance, Reputation, and CEO Pay. *Administrative Science Quarterly*, *45*(3), 523-556. Cornell University Samuel Curtis Johnson Graduate School of Management.
- Stogdill, R. M. (1974). *Handbook of leadership: A survey of theory and research*. New York: Free Press.
- Strenta, A. C., & Elliott, R. (1987). Differential Grading Standards Revisited. *Journal of Educational Measurement*, *24*(4), 281-291. doi:10.1111/j.1745-3984.1987.tb00280.x
- Timmons, J. A. (1978). Characteristics and Role Demands of Entrepreneurship. *American Journal of Small Business*, *3*(1), 5-17.

Windschitl, P. D., Kruger, J., & Simms, E. N. (2003). The influence of egocentrism and focalism on people's optimism in competitions: when what affects us equally affects me more.

Journal of Personality and Social Psychology, 85(3), 389-408. doi:10.1037/0022-3514.85.3.389

Wolfers, J. (2002). Are Voters Rational? Evidence from Gubernatorial Elections.

Woo, E. (1997). Boalt Hall Law School Kills Its Grade-Weighting Policy. *Los Angeles Times*.

Los Angeles, CA.

Tables

Table 1

Summary of ANOVA results by dependent variable and factor (Study 1)

	Situation (Avg. GPA)			Performance (GPA)			Avg GPA * GPA		
	F (2, 44)	p	η_p^2	F (2, 44)	p	η_p^2	F (4, 88)	p	η_p^2
Success	61.75	< .001	0.74	48.02	< .001	0.69	3.42	0.012	0.13
Prob. of Accept	58.97	< .001	0.79	52.79	< .001	0.76	7.74	< .001	0.62
Admission	42.43	< .001	0.66	52.79	< .001	0.76	3.42	< .001	0.23
Mean Z-score	94.24	< .001	0.81	85.09	< .001	0.80	5.88	< .001	0.21

Table 2

Probability of being accepted by condition (Study 1).

		GPA (Performance)			
		Avg. - 0.3	Avg.	Avg. + 0.3	Total
Average GPA (Situation)	Low (2.4)	.00	.09	.26	.12
	Med (3.0)	.04	.57	.87	.49
	High (3.6)	.30	.91	.96	.72
	Total	.12	.52	.70	

Table 3

Ratings of performance and promotion worthiness based on candidate performance and airport difficulty (Study 2). Average ratings of the candidates appear in each cell. Standard deviations appear in parentheses.

Situation (Historic on-time %)	Easy \approx 85%			Difficult \approx 70%		
	Performance (Relative on-time %) Low (5% below)	Medium (average)	High (5% above)	Low (5% below)	Medium (average)	High (5% above)
Performance Rating	3.54 (1.10)	4.79 (0.85)	6.14 (0.91)	3.09 (1.03)	4.25 (0.80)	5.55 (0.94)
Promotion Worthiness Rating	3.18 (1.23)	4.30 (1.07)	5.84 (0.96)	2.71 (1.16)	3.76 (1.10)	5.19 (1.08)

Table 4

Applicant level descriptive statistics and bivariate correlations (Study 3); all correlations are statistically significant ($p < .001$)

	Mean	S.D.	1	2	3	4	5	6	7
1. Male	0.72	0.45							
2. Age	29.8	4.56	.15						
3. US Citizen	0.53	0.50	-.07	-.07					
4. GMAT	669	65.4	.14	-.16	-.08				
5. Interview rating	0.00	0.96	-.05	.05	.05	.05			
6. Years of work history	5.35	3.16	.08	.39	-.09	.05	.07		
7. Undergraduate GPA	3.30	0.39	-.15	-.21	.03	.16	.04	-.13	
8. Admitted	0.31	0.46	-.04	.03	.16	.17	.38	-.02	.14

Table 5

Undergraduate institution level descriptive statistics and bivariate correlations (Study 3); all correlations are statistically significant ($p < .05$) unless marked ()*

	Mean	S.D.	1	2	3	4	5	6
1. Gourman quality score	3.5	0.6						
2. US News quality score	61.7	16.1	.31					
3. Average entrance SAT	1153	144	.51	.81				
4. Average GPA	3.2	0.2	.18	.62	.51			
5. Tuition	18450	9686	.12	.66	.63	.51		
6. University	0.6	0.5	.38	-.21	-.18	-.17	.02*	
7. Private	0.5	0.5	-.13	.48	.40	.50	.82	-.14

Table 6.

Summary of ordinal logistic models predicting candidates' admission outcomes (Study 3).

	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6	Model 7
Deny Waitlist	0.571*** (0.016)	6.860*** (0.045)	8.500*** (0.022)	10.359*** (0.037)	7.897*** (0.026)	8.305*** (0.050)	9.173*** (0.024)
Waitlist Accept	0.840*** (0.016)	7.370*** (0.049)	9.027*** (0.028)	10.913*** (0.042)	8.448*** (0.032)	8.884*** (0.055)	9.544*** (0.027)
Performance: Relative GPA	0.868*** (0.042)	0.771*** (0.068)	0.733*** (0.073)	0.798*** (0.075)	0.972*** (0.078)	1.057*** (0.081)	1.562*** (0.058)
Situation: Average GPA	1.389*** (0.097)	0.359* (0.179)	0.250 (0.197)	0.323 (0.203)	0.736*** (0.203)	0.878*** (0.208)	1.597*** (0.154)
Interview		1.012*** (0.030)	1.021*** (0.030)	1.077*** (0.031)	1.025*** (0.031)	1.080*** (0.032)	
Years of work experience		-0.002 (0.009)	-0.004 (0.009)	-0.075*** (0.012)	0.029** (0.011)	0.026 (0.015)	0.123*** (0.012)
GMAT		0.010*** (0.000)	0.013*** (0.000)	0.014*** (0.000)	0.014*** (0.000)	0.015*** (0.000)	0.016*** (0.000)
School Quality: US News Rating		0.015*** (0.002)	0.011*** (0.002)	0.011*** (0.002)	0.015*** (0.002)	0.014*** (0.002)	0.018*** (0.002)
Age			0.003 (0.006)	0.035*** (0.008)	0.006 (0.008)	-0.024* (0.011)	-0.061*** (0.008)
Gender: Male			-0.614*** (0.056)	-0.648*** (0.057)	-0.696*** (0.057)	-0.714*** (0.059)	-0.560*** (0.041)
Citizenship: US			0.213* (0.098)	0.241* (0.101)	0.248* (0.100)	0.240* (0.104)	0.086 (0.074)
Race: African American			1.266*** (0.144)	1.406*** (0.151)	1.496*** (0.148)	1.643*** (0.154)	1.547*** (0.107)
Race: Asian			-0.154* (0.068)	-0.174* (0.070)	-0.205** (0.070)	-0.214** (0.072)	-0.191*** (0.050)
Race: Hispanic			1.229*** (0.152)	1.280*** (0.158)	1.374*** (0.155)	1.422*** (0.162)	0.943*** (0.101)
Race: American Indian			0.523*** (0.001)	0.474*** (0.003)	0.640*** (0.002)	0.655*** (0.005)	0.648*** (0.004)
Race: Other			-0.121 (0.070)	-0.143* (0.073)	-0.126 (0.072)	-0.122 (0.075)	-0.146** (0.055)
Application Years				Included		Included	Included
Admitting School					Included	Included	Included
N	19503	8681	8674	8674	8674	8674	17504
AIC	31774	13323	13032	12598	12627	12208	22235

* p < .05; ** p < .01; *** p < .001

Standard errors are in parentheses.

Table 7

Summary of ordinal logistic models of graduate school admissions, comparing standardized measures of applicants' undergraduate institution quality (Study 3).

	Model 8 B (SE)	Model 9 B (SE)	Model 10 B (SE)	Model 11 B (SE)
Deny Waitlist	7.847*** (0.023)	8.050*** (0.022)	7.829*** (0.023)	7.656*** (0.024)
Waitlist Accept	8.218*** (0.026)	8.419*** (0.025)	8.201*** (0.026)	8.026*** (0.027)
Performance: Relative GPA	1.562*** (0.059)	1.419*** (0.055)	1.529*** (0.057)	1.584*** (0.060)
Situation: Average GPA	1.597*** (0.156)	2.274*** (0.130)	1.627*** (0.143)	1.556*** (0.158)
Years of work experience	0.123*** (0.012)	0.132*** (0.011)	0.125*** (0.011)	0.126*** (0.012)
GMAT	0.016*** (0.000)	0.017*** (0.000)	0.016*** (0.000)	0.016*** (0.000)
Standardized School Quality: US News	0.294*** (0.026)			0.261*** (0.038)
Standardized School Quality: Gourman		0.055* (0.023)		-0.055* (0.027)
Standardized School Quality: SAT			0.280*** (0.026)	0.085* (0.039)
Age	-0.061*** (0.008)	-0.072*** (0.008)	-0.064*** (0.008)	-0.064*** (0.009)
Gender: Male	-0.560*** (0.041)	-0.605*** (0.040)	-0.554*** (0.041)	-0.543*** (0.042)
Citizenship: US	0.086 (0.074)	0.133 (0.071)	0.090 (0.072)	0.044 (0.074)
Race: African American	1.547*** (0.107)	1.605*** (0.103)	1.577*** (0.105)	1.500*** (0.111)
Race: Asian	-0.191*** (0.050)	-0.204*** (0.049)	-0.212*** (0.049)	-0.181*** (0.051)
Race: Hispanic	0.943*** (0.101)	0.965*** (0.098)	0.921*** (0.099)	0.936*** (0.102)
Race: American Indian	0.648*** (0.004)	0.802*** (0.005)	0.726*** (0.004)	0.565*** (0.003)
Race: Other	-0.146** (0.055)	-0.145** (0.053)	-0.156** (0.054)	-0.153** (0.056)
Application Years	Included	Included	Included	Included
Admitting School	Included	Included	Included	Included
N	17504	18441	18121	16926
AIC	22235	23632	23123	21507

* p < .05; ** p < .01; *** p < .001

Standard errors are in parentheses.

Table 8

Summary of alternative analysis approaches (Study 3).

	Model 7 B (SE)	Model 12 B (SE)	Model 13 B (SE)	Model 14 B (SE)
Deny Waitlist	9.173*** (0.024)	9.155*** (0.024)	8.333*** (0.040)	14.244*** (0.016)
Waitlist Accept	9.544*** (0.027)	9.527*** (0.027)	8.691*** (0.043)	14.616*** (0.020)
Performance: Relative GPA	1.562*** (0.058)	1.568*** (0.058)	1.685*** (0.075)	
Performance: Nominal GPA				1.562*** (0.051)
Situation: Average GPA	1.597*** (0.154)	1.397*** (0.148)	1.715*** (0.217)	0.035 (0.146)
Relative GPA x Average GPA		1.250*** (0.032)		
Years of work experience	0.123*** (0.012)	0.125*** (0.012)	0.130*** (0.015)	0.123*** (0.012)
GMAT	0.016*** (0.000)	0.016*** (0.000)	0.015*** (0.000)	0.016*** (0.000)
School Quality: US News Rating	0.018*** (0.002)	0.018*** (0.002)	0.017*** (0.003)	0.018*** (0.002)
Age	-0.061*** (0.008)	-0.062*** (0.008)	-0.081*** (0.011)	-0.061*** (0.009)
Gender: Male	-0.560*** (0.041)	-0.556*** (0.041)	-0.511*** (0.049)	-0.560*** (0.042)
Citizenship: US	0.086 (0.074)	0.078 (0.074)	-0.154 (0.092)	0.086 (0.074)
Race: African American	1.547*** (0.107)	1.567*** (0.107)	1.439*** (0.134)	1.547*** (0.106)
Race: Asian	-0.191*** (0.050)	-0.190*** (0.050)	-0.181** (0.058)	-0.191*** (0.050)
Race: Hispanic	0.943*** (0.101)	0.953*** (0.101)	0.986*** (0.124)	0.943*** (0.101)
Race: American Indian	0.648*** (0.004)	0.669*** (0.004)	0.697*** (0.004)	0.648*** (0.003)
Race: Other	-0.146** (0.055)	-0.143** (0.055)	-0.078 (0.066)	-0.146** (0.055)
Application Years	Included	Included	Included	Included
Admitting School	Included	Included	Included	Included
N	17504	17504	11361	17504
AIC	22235	22219	15018	22235

* p < .05; ** p < .01; *** p < .001

Standard errors are in parentheses.

Figure 1. Model 7 predicted probability of acceptance by situation and performance (Study 3)

