First Impression Biases and the Value of Blind Auditioning

Jasmin Droege
St Cross College
University of Oxford

A thesis submitted for the degree of
Master of Philosophy in Economics

Trinity 2019

Wordcount: 20,088

To Matthew.

Acknowledgements

I would like to thank Dr Willemien Kets for her encouragement and guidance as my supervisor. I would like to thank Dr Matthew J. Robertson for the inspiration to pursue this line of research, the countless proofreading marathons and the twenty-four seven LaTeX emergency support. I would like to thank Dr Elizabeth Baldwin for her encouragement as my departmental supervisor during the time leading up to this line of research. Finally, I would like to thank my Parents and my Twin Sister for all their support and faith in me.

Abstract

Motivated by the performing arts, I investigate the repercussions of an evaluator's bias against a specific group of applicants, such as women, in a setting where the evaluator decides upfront between holding an informed or a blind audition. In the latter case, the evaluator learns neither the applicant's ability nor the gender. I show that a blind audition comes with a fundamental trade-off: it avoids misleading first impressions but prohibits the screening of applicants by ability. Above a threshold bias, the evaluator prefers a blind audition to provide high effort incentives exclusively for highly able applicants while low-ability applicants do not participate. For a low bias, requiring supplementary information about the applicant's type acts as an efficient screening device and, thereby, provides such targeted effort incentives. From a policy perspective, committing to no information can, therefore, be beneficial for the evaluator if he knows that he would otherwise not be impartial. Applicants' audition preferences differ by ability if the evaluator's bias against female applicants is low and differ by gender if the bias is large. Surprisingly, the preferences of a highly biased evaluator align with those of a highly able female. My model explains the empirical finding that blind auditions have increased the probability of women being hired via taste-based discrimination and challenges competing explanations grounded in statistical discrimination. Additionally, I show that uncertainty introduces the sizeable risk of market failure and may render informed auditions more profitable even if the evaluator is highly biased. From a policy perspective, ability-targeting interventions under uncertainty are, therefore, crucial to ensure equality of opportunity and to avoid zero-hour contracts.

Keywords: first impression, cognitive bias, blind audition, taste-based discrimination.

Contents

List of Figures

1

Introduction

> *"I've been in auditions without screens, and I can assure you that I was prejudiced. I began to listen with my eyes, and there is no way that your eyes don't affect your judgment. The only true way to listen is with your ears and your heart."*
> – Julie Landsman in Gladwell, *Blink: The Power of Thinking Without Thinking*

In the performing arts, applicants participate in a form of audition to secure a job. In the design of this audition, the evaluator, who is responsible for hiring the the most talented candidate, commits to one of two mechanisms; that is, a blind or informed audition. If informed, the evaluator requires applicants to submit further information before audition day. This often takes the form of a CV that contains valuable information about the applicant's ability. One proxy is the school an applicant went to, such as Juilliard, which increases the evaluator's belief about the applicant's ability. A CV, however, also contains irrelevant information about gender or race which may lead to bias of the evaluator; that is, misleading first impressions that have a *discounting effect* on more relevant information obtained at a later stage (Thorngate et al., 2010). If blind, the evaluator receives no further information about the applicant before audition day. What is more, as the evaluator cannot see the applicant during her performance, the hiring decision is only based on sound. Given the form of audition, the applicant decides upon how much effort to invest in preparing a performance which can manifest in the difficulty of the piece performed on audition day. Importantly, this sequential decision-making of the evaluator and the applicant allows for systematic effects of the audition form on effort incentives.

In this thesis, I analyse the repercussions of the evaluator's bias against a specific group of applicants. I ask the question whether, and under which conditions, it can be rational for the evaluator to ignore supplementary information about an applicant to avoid misleading first impressions. Moreover, I ask whether a move to a blind audition can lead to more talented candidates being hired; in essence, whether, and under which conditions, a blind audition provides better effort incentives exclusively for highly able applicants. I focus on a bias against female applicants and frequently use the context of an orchestra to build intuition. However, my research questions are pertinent to a broader range of contexts in which the evaluator requires the applicant to prepare a project for review, and possibly submit additional information susceptible to bias, before deciding upon acceptance.

My thesis is inspired by the observation that many U.S. symphony orchestras moved to blind auditions in the 1970s and 1980s (Goldin and Rouse, 2000). In particular, the use of curtains has been marketed as a fairer form of audition and equal opportunity policy, protecting women and minorities against discrimination (Karl Schiebler in The The Economist, 1996). The practice of blind auditions is often credited to have its origins in 1970 with the request of the late African American bassist Art Davis for candidates to play behind a curtain when auditioning for the New York Philharmonic. When the Philharmonic refused to conceal applicants' identities, Davis took the orchestra to court for racial discrimination.

The increasing focus on equal opportunity in the performing arts has also spawned more unconventional entertainment formats. The singing competition franchise "The Voice", adapted in 129 regions across the world, uses blind auditions in the first round. Therefore, according to creator John de Mol, the judges focus on quality alone and identify genuine singing talent (Universal Music Group, 2011). When the judges listen, they face the audience rather than the contestant. Only if a judge commits to taking on the contestant, can she press a button which turns her chair around.

Blind hiring has the potential to be implemented in industries other than the performing arts. This depends on whether interviews and CVs to review candidates can be substituted or anonymised. The technology company GapJumpers offers blind hiring to firms in preliminary rounds. The Silicon Valley start up has developed software that withholds information about ethnicity, gender, age or educational background[1]. It also requires the applicant to complete an online test, tailored to the skills in the job specification. This allows candidates to essentially blind-audition for a job, and employers to signal that "[i]f you can show me your skills in this role, I am willing to interview you, regardless of where you come from, what you look like or who you are" (CEO Kedar Iyer in Cain Miller, 2016). Academia may be another application if blind auditioning can help to address the lower success rates of female applicants for funding schemes such as ERC grants (Vernos, 2013).

I address the two polarised views on the use of blind auditions to hire the most talented candidate. Proponents like Gladwell (2005) argue that blind auditions have helped evaluators avoid misleading first impressions, what he calls "snap judgments". He bases

[1]The act of withholding applicants' names may already have a significant effect. See, for example, Bertrand and Mullainathan (2004) for an experiment in which CVs are randomly assigned African-American- or White-sounding names.

his argument on the prejudices that women have encountered in the past when auditioning for orchestra positions. Two examples of such are not possessing the physical strength for instruments traditionally considered "male", or an inferior attitude and resilience compared to males. Opponents fear that, while avoiding bias, the evaluator also loses valuable information that can screen applicants by ability. As a result, top talent may be lost in an increasingly large applicant pool and standards of fairness may actually deteriorate (Holland, 1981).

My contribution is to show that these two polarised views are not mutually exclusive. I argue that a blind audition comes with a fundamental trade-off: it grants impartiality that protects women, blacks and potentially other groups from the evaluator's bias, but prohibits the screening of applicants by ability. I provide conditions on the severity of the bias under which the evaluator prefers a blind audition to provide high effort incentives exclusively for highly able applicants while low-ability applicants do not participate and, thereby, show that the two polarised views are actually not contradictory. A second contribution is to provide an explanation for the reduced form findings of Goldin and Rouse (2000); that is, the use of curtains is associated with a 50 percent increase in the probability that a female will be advanced from preliminary rounds. The authors also find a severalfold increase in the probability that a female will be hired when curtains are used in the audition.

My results on the evaluator's audition preferences are threefold. First, there is a threshold bias above which the evaluator prefers a blind audition to provide *targeted* effort incentives: high effort incentives exclusively for highly able applicants while low-ability applicants do not participate. Conversely, in an informed audition the evaluator's bias would *distort* effort incentives: males, irrespective of their ability, would exert low effort, resting on their laurels and exploiting their ability to identify themselves as male. Highly able females, on the other hand, would be deterred from exerting high effort because the benefit would be very small relative to the higher effort costs that they would have to bear. Finally, low-ability females would drop out of the audition. From a policy perspective, this first result highlights that a strategy of commitment to no information can be beneficial if the evaluator knows that he would otherwise not be impartial. It also highlights that equal opportunity and targeted effort incentives are complementary rather than conflicting objectives above a certain threshold bias. In other words, there is no trade-off between implementing blind auditions as a policy to counteract first impression biases and efficiency due to adverse effects on targeted effort provision. If anything, implementation of equal opportunity is effort enhancing for the highly able applicants above this threshold.

Second, for a low bias, requiring a CV acts as an efficient screening device and, thereby, provides targeted effort incentives. Intuitively, the evaluator is able to focus on the relevant information without being sidetracked by the noise on the CV. He can, therefore, use the CV as a means to attract a pool of exclusively high-ability applicants. In a blind audition, in contrast, the evaluator would also attract applicants of low ability. Furthermore, the evaluator would have to give up the advantage of discriminating between applicants of different characteristics; that is, when all exert high effort, the evaluator would have no means of distinguishing them behind the curtain. Therefore, he would need to hire all applicants with the same probability. From a policy perspective, this second result highlight that, for a low bias, equal opportunity and targeted effort incentives are conflicting objectives: while blind auditions serve as a policy to counteract first impression biases, they may lead to an avalanche of auditionees and efficiency loss as unintended consequences.

Third, for an intermediate range of bias, the evaluator's audition preferences depend on the characteristics of the applicant pool. If predominantly of low ability, there are significant losses from hiring low ability males in an informed audition relatively frequently. If predominantly of high ability, these losses are secondary and it is optimal to hire with differential probabilities.

My results on applicants' audition preferences are twofold. First, if the evaluator's bias against female applicants is low, preferences differ by ability. This is because the highly able prefer to identify themselves as such, whereas the less able would prefer to pool with the highly able. Conversely, if the evaluator's bias against female applicants is large, preferences differ by gender. Intuitively, the bias plays out in the males' favour and, therefore, they prefer to identify themselves in an informed audition. Highly able females, on the other hand, rather trade off not being able to reveal their high ability for not having to reveal their gender. As a corollary, the evaluator's preferences align with the preferences of high-ability females if his bias against them is large. This is surprising: one would expect their preferences to be diametrically opposed if the evaluator is not impartial. Second, my model sheds light on the underlying forces of the findings of Goldin and Rouse (2000), and shows that the gains from being able to conceal one's gender need not be uniform. For a low bias, the authors' observations are driven by the possibility of hiring more low-ability females in the move to a blind audition. In fact, this more than compensates for the reduced probability of hiring high-ability females. Conversely, if the evaluator's bias is large, the authors' observations are entirely driven by hiring high-ability females more often in the move to a blind audition, with no change for their low-ability counterparts.

My results on the introduction of asymmetric uncertainty are twofold. First, if there is a non-negligible degree of uncertainty, the expected net utility gain for a highly biased evaluator from moving to a blind audition is smaller. In fact, a blind audition is no longer guaranteed to be more profitable for him. This is because low-ability applicants are no longer deterred from participating in a blind audition if the evaluator is highly biased. The slightest chance of attaining high performance quality as a consequence of randomness or having a lucky day makes participation attractive for the less able. As a result, the losses from low-ability-low-effort applicants may outweigh the benefits from high-ability-high-effort applicants in a blind audition. Second, there is the sizeable risk of market failure under uncertainty: for certain bias-prior-uncertainty combinations, neither blind nor informed auditions are profitable and, therefore, the evaluator might not want to hold an audition at all. Consequently, uncertainty does not lead to a redistribution of the evaluator's preferences. Rather, he is worse off because uncertainty reduces the attractiveness of a blind audition by more than it increases the attractiveness of an informed audition. From a policy perspective, ability-targeting interventions are key if uncertainty is an inherent feature of the environment. First, it guarantees the profitability of a blind audition when the evaluator is highly biased and, thus, avoids market failure in the form of continual subcontracting or zero-hour contracts. Second, it guarantees that blind auditions still do better than informed auditions. This may be desirable for equality of opportunity along the gender dimension of applicants.

1.1 Literature

This thesis relates to several strands of literature. First, I draw upon the literature on cognitive biases. In particular, I argue that the evaluator in an informed audition is likely subject to different types of cognitive biases which may be summarised in a reduced form bias parameter. One potential driver is intergroup bias; that is, the evaluator may judge members of his own group, such as gender or race, more favourably. Rainer Kuechl from the Vienna Philharmonic, as an instance of the former, claims that "women play differently - not worse but more softly, more flexibly" (Chadwick, 1997, pp.1324-1326). An instance of the latter are significant differences in a panel's grade predictions by race when the auditionee performs a classical excerpt (Clauhs, May 2013). A second potential driver of bias is race and gender associations with instruments, as mentioned by Gladwell (2005). Flute or harp, for example, are traditionally described as female instruments. While gender stereotyping has decreased, these associations persist due to parents, peer pressure and music teachers (Abeles, 2009). A third potential driver motivating the modelling of bias is interaction effects of gender and race in music evaluation. Early experiments, for instance, showed that black males scored lowest among black students and white females scored lowest among white students (Elliott, 1995). Lastly, there are other non-musical factors in performance evaluation that warrant the modelling of bias. This could be as subtle as the strength of the applicant's handshake at the audition if this opening ritual, in turn, influences the evaluator's belief about the applicant's strength of character. In particular, "[b]ecause men tend to have greater upper body strength than do women, a male's handshake is likely to be firmer than a female's handshake" and, as a result, "if the [evaluator] shakes a woman's hand, it is more likely the shake will fail the [evaluator]'s character test", leading him to discount more relevant information obtained from the female's subsequent performance (Thorngate et al., 2010, p.51). Other examples are an attractiveness bias or the stage presence of an applicant, manifesting in non-verbal behaviour, gestures, facial expressions or attire. In particular, there is evidence for a vision heuristic in informed auditions. Yet, evaluators consistently report to value sound as central in performance. As a result, they arrive at different winners depending on whether they have visual information about the applicant or not. Therefore, "professional judgment appears to be made with little conscious awareness that visual cues factor so heavily in preferences and decisions" (Tsay, 2013, p.14583).

Moreover, I draw upon contract theory by assuming verifiable disclosure: in an informed audition the applicant cannot misrepresent her gender nor ability on her CV (Milgrom, 1981; Grossman, 1981). Conversely, the evaluator cannot commit to cherry-pick information relevant to him, such as in the form of technology offered by GapJumpers. Effort incentives are, therefore, influenced by the evaluator's bias as it enters his hiring rule. This feature is similar to implicit incentives arising in a dynamic setting (Meyer and Vickers, 1997). Second, the evaluator has a menu of audition forms available to hire the most talented candidate and, given his commonly known bias, needs to choose the optimal form (Laffont and Tirole, 1986). Third, the blind audition is akin to a signalling model à la Spence (1973) as the applicant's effort decision may be informative about the ability dimension of her type while gender cannot be signalled. Therefore, I study both pooling and separating equilibrium in this form of audition.

Third, this thesis builds on literature in the economics of discrimination. In particular, one theoretical explanation for the findings of Goldin and Rouse (2000) is a model of statistical discrimination (Taylor and Yildirim, 2011). The authors endow the applicant with a one-dimensional type that reflects her true productivity. This type can be interpreted as innate ability or experience. Therefore, the submission of a CV in an informed audition implies a perfectly informative signal about ability. The evaluator essentially obtains just the right amount of information to accurately judge the applicant's ability. However, Taylor and Yildirim (2011) mention "the possibility of psychological bias on the part of the evaluator and the role of blind [auditions] in mitigating prejudice" (p.786) as an intriguing avenue for future research.

This thesis fills the gap in the literature: I argue that, in reality, a CV is a composite of valuable information about ability and noise. This may trigger misleading first impressions;

that is, "snap judgments" on part of the evaluator that may prevent him from hiring the most talented candidate. Therefore, I provide an alternative model grounded in taste-based discrimination à la Becker (1971). I endow the applicant with a two-dimensional type that can be interpreted as a categorisation of information as valuable, like ability, or irrelevant information about the applicant, like gender, for inferring ability. I proceed with the assumption that, for the same level of ability, females and males are equally productive. First, this assumption is in line with the notion that "[m]usic expresses life and life's emotions, but has no gender" (Starr, 1974, p.14) and cases such as Abbie Conant (Gladwell, 2005, p.245-248) underline the empirical validity of this assumption[2]. Second, this assumption stresses that prejudices against female musicians, social attitudes and self-identity rather than physical limitations are key in explaining the findings of Goldin and Rouse (2000)[3].

The remainder of this thesis is structured as follows. In Chapter 2, I set out preliminaries common to both audition forms; that is, timing, preferences, simplifying assumptions regarding the performance outcome and solution concepts. In Chapter 3, I solve for, and characterise, equilibrium in the informed audition for all levels of bias. In Chapter 4, I solve for, and characterise, pooling and separating perfect Bayesian equilibrium in the blind audition for all levels of bias. In Chapter 5, I offer a comparison of the applicant's effort in the blind and informed audition for all levels of bias, and characterise the hiring process that gives the evaluator the highest ex-ante expected net utility. I continue with a discussion of the applicant's preferences over auditions for different types and levels of bias. I conclude this chapter with a comparison to the main competing statistical discrimination model. In Chapter 6, I introduce asymmetric, empirically motivated uncertainty into the audition process to challenge the competing statistical discrimination model also in an uncertain environment and allow for richer policy recommendations. Chapter 7 concludes with avenues for future research. Appendix A contains a version of the model under certainty that features multiplicative technology, in which I show that my main results are robust to a specification in which being of high ability is more valuable to the evaluator. In Appendix B, I exhaustively show the robustness of the evaluator's beliefs in a separating blind audition.

[2] In the case of Abbie Conant, medical tests, such as blowing through special machines to measure Conant's lung capacity, blood samples to measure her capacity for absorbing oxygen or a chest exam, confirmed that a female can possess the necessary physical strength for the solo trombone (Gladwell, 2005, p.245-248).

[3] See Seltzer (1989, pp.211-221) for a discussion of the common prejudices against female musicians. Starr (1974) lists the common myths about the limitations of female performers. Being a renowned solo pianist, she takes her profession as an example, and argues that "there are many male pianists who sound smaller-scaled and more effeminate than the great female virtuosi, who have *no physical limitations*, and can play any piece of music written for the instrument. They can achieve as much power, strength and stamina as any male, though they may have to use more energy" (p.14).

2

Preliminaries

2.1 Timing

Two risk neutral parties, an applicant (the agent) and an evaluator (the principal) play a three-stage game; that is, $s \in S = \{1, 2, 3\}$.

In the first stage, the evaluator commits to a hiring process, which is either a blind or an informed audition, as well as an exogenously determined wage w. If the audition is informed, the evaluator learns the applicant's type $\theta = (\eta, g)$ where $\eta \in \{\eta_L, \eta_H\} \equiv \{1, 2\}$ is ability and $g \in \{m, f\} \equiv \{0, 1\}$ is gender. If the audition is blind, the evaluator does not learn the applicant's type. He only knows the type distribution; that is, he knows that the applicant is female or male with equal probability. The applicant is of high ability with $\Pr(n_H) = p \in (0, 1)$. It is common knowledge that ability and gender are independent.

In the second stage, the applicant moves, knowing her type and whether it is a blind or informed audition. She chooses how much effort to invest in preparing a performance, where effort $e \in \{e_L, e_H\} \equiv \{1, 2\}$ can be low or high. There is an outside option O for the applicant that yields a zero payoff.

In the third stage, the applicant delivers her performance. It can turn out to be of low or high quality; that is, $q \in \{q_L, q_H\} \equiv \{1, 2\}$. The evaluator makes the hiring decision based on type θ (i.e. ability and gender) and performance quality q if the audition is informed, and based on performance quality q alone if the audition is blind. In either form of audition, the evaluator has a random outside option $\bar{U}$, which he can revert to instead of hiring the applicant.

2.2 Preferences of Players

The risk-neutral evaluator is assumed to be biased against female applicants. This is represented in form of a bias parameter $\beta \in [0, 2]$ which is common knowledge in either form of audition. I formalise the evaluator's preferences with a gross utility function $V(q, \eta, g)$

consisting of two parts. The first part $f(q, \eta) = q + \eta$ is a production function according to which the applicant contributes to the orchestra if hired[1]. The evaluator, therefore, prefers to hire an applicant who is of high ability and delivers a high-quality performance in stage 3. In particular, given the applicant's ability η, which the evaluator may learn in stage 1, and her performance quality q, always observed in stage 3, the evaluator earns revenue $f(q, \eta)$ by hiring the applicant. This assumption may be justified if $f(q, \eta)$ is regarded more broadly as an applicant's marketable talent, consisting of a constant and time-variant component[2]. Innate ability as the former can manifest in playing an unseen piece to a satisfactory standard. Effort as the latter can manifest in improved playing skills and being able to perform the piece that was mastered in preparation for the audition as a solo. Note that this formulation also implies that males and females are perfect substitutes in production as gender does not enter $f(q, \eta)$; that is, for the same level of ability and performance quality, males and females are equally productive. The second part βg captures the evaluator's bias against female applicants[3]:

$$V(q, \eta, g) = f(q, \eta) - \beta g = q + \eta - \beta g \tag{2.1}$$

Importantly, the gross utility function specified in (2.1) does not include wages. This is because w is assumed to be the same for each applicant and the outside option, and is exogenously determined[4].

Assumption 1. $w \equiv \mathbb{E}[\bar{U}] = w_g = w_{\bar{U}}$ *for* $g \in \{m, f\}$.

I can, therefore, abstract from wages in the evaluator's utility function. Wage costs are sunk and do not affect his hiring decision in stage 3. Later, I will refer to the evaluator's utility net of wage costs as Π.

Assumption 2. $\bar{U} \sim U[2 - \beta, 4]$.

[1]Under a multiplicative form of the production function, $f(q, \eta) = q\eta$, most of the results remain qualitatively unchanged. See Appendix A for a discussion of the key differences.

[2]See, for example, Tsay and Banaji (2011) for the two sources of talent and how they are perceived by expert decision-makers. The authors characterise "naturals" via early evidence of high innate ability, while "strivers" are characterised by high motivation and perseverance.

[3]This form of utility function is inspired by Becker (1971). He presents a model of taste-based discrimination (chapter 3, p.39) where an employer has a taste parameter d ("coefficient of discrimination") in his utility function so that the wage of a member of the minority group N is effectively $w_N + d$.

[4]This assumption is in line with orchestras' practice to specify a base pay according to which positions are remunerated, making wages independent of the audition form. One example in the UK are employment agreements negotiated by the Musicians' Union (n.d.), which are applicable to members employed with mayor orchestras like BSO, BBC or RSNO. Besides having empirical support, this assumption eliminates the screening aspect that wage setting would otherwise have in stage 1. In particular, endogenising wages would make a blind audition preferable even if the evaluator was not biased due to significant wage savings. Furthermore, the evaluator would then be best off setting $w = 0$ and always hiring the outside option because $\bar{U}$ does not respond to incentives in this model.

The evaluator has a random outside option $\bar{U} \sim U[2-\beta, 4]$ which he can revert to instead of hiring the applicant. The value of $\bar{U}$ is unknown to the evaluator when making the hiring decision in either form of audition. This assumption may be justified if $\bar{U}$ is thought of as the possibility to engage a substitute player from an agency he can contract with. This agency randomly provides the same four types of players who then choose effort $e \in \{e_L, e_H\}$ nonstrategically. A substitute player, therefore, generates gross utility of at least $2 - \beta$ if she is of type $\theta = (\eta_L, f)$ and exerts low effort. A substitute player generates at most 4 if he is of type $\theta = (\eta_H, m)$ and exerts high effort. Because the evaluator cannot contract upon type and effort choice of the substitute player, $\bar{U}$ is essentially a gamble that may turn out to be better or worse than the applicant.

The evaluator can face four possible types of applicants in this game given the assumptions on η and g, $\Theta = \{(\eta_L, m), (\eta_H, m), (\eta_L, f), (\eta_H, f)\}$. The applicant, whether of low or high ability, wants to be hired to receive a a wage w net of effort costs. In particular, by choosing $e \in \{\eta_L, \eta_H\}$ she maximises her expected payoff, taking into account how her effort choice influences the probability of being hired via its effect on performance quality. Her effort costs are given by $c(e) = \frac{e^2}{2\eta}$, which is decreasing in her ability. An applicant receives a zero payoff if she is not hired and her outside option gives zero utility. These assumptions allow me to simplify the applicant's expected utility in an informed audition to:

$$U(e, \eta, g) = \Pr(\text{hired}|q, \eta, g)w - \frac{e^2}{2\eta} \tag{2.2}$$

In a blind audition, the applicant's form of utility function remains unchanged except for the hiring probability never directly depending on gender or ability; that is, $\Pr(\text{hired}|\cdot)$ changes to $\Pr(\text{hired}|q)$.

2.3 Performance Outcome and Solution Concept

In the first part of the thesis, I focus on the interaction of the strength of the evaluator's bias β and the information structure of the audition. I abstract from any moral hazard considerations; that is, effort maps one-to-one into performance quality, and can be inferred by the evaluator. In Chapter 6, I relax this assumption and allow for moral hazard by considering settings where low effort does not uniquely determine performance quality.

Assumption 3.

$$\Pr(q_H|e_H) = 1 \textit{ with } \Pr(q_H|e_L) = 0$$

and

$$\Pr(q_H|e_L) = 0 \textit{ with } \Pr(q_L|e_L) = 1.$$

Given Assumption 3, the performance will be of high quality in stage 3 if the applicant exerts high effort in stage 2. Conversely, it will be of low quality if the applicant exerts low effort. Therefore, an informed audition is a game of complete information and the solution concept is backward induction. Having specified a hiring rule in stage 3, I solve for an applicant's optimal effort level for different values of β in stage 2. Lastly, I solve for the evaluator's expected utility for different values of β in stage 1.

A blind audition is a game of incomplete information and the solution concept is perfect Bayesian equilibrium, refined by the D1-Criterion when necessary. I specify a belief for the evaluator at his information sets q_L and q_H in stage 3. Given these beliefs and the hiring rule, I solve for an applicant's optimal effort level for different values of β in stage 2 and the evaluator's expected utility for different values of β in stage 1.

3

Informed Audition

If the evaluator has decided for an informed audition and committed to a wage $w \equiv \mathbb{E}[\bar{U}]$, stage 2 and 3 can be viewed as a game of complete information with four proper subgames; that is, one subgame for each applicant type. I, therefore, solve the game by first considering the evaluator's hiring decision in $s = 3$. Using this decision, I then determine the applicant's optimal effort choice in $s = 2$ (see Figure 3.1).

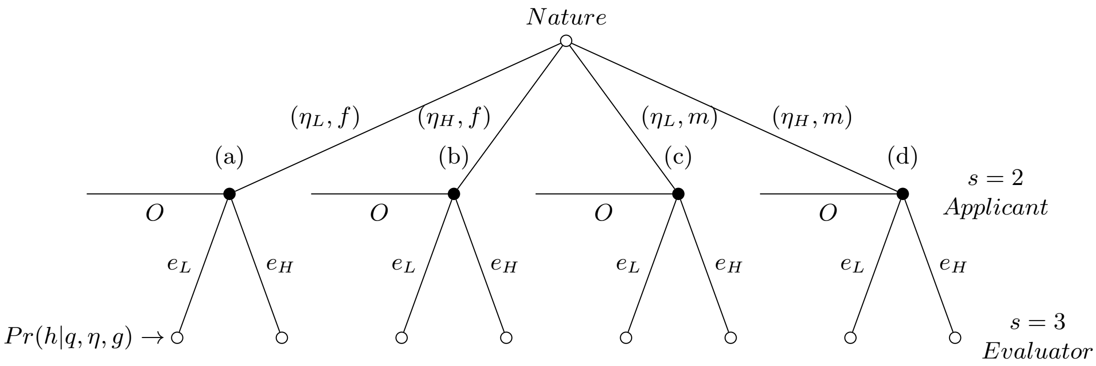

Figure 3.1: Four Proper Subgames in Informed Audition[1]

3.1 Hiring Decision in Stage 3

Due to the random outside option $\bar{U}$, the evaluator hires the applicant in stage 3 with a probability derived from comparing the utility from hiring an applicant to the cumulative distribution function of $\bar{U}$:

$$\text{Pr}(\text{hired}|q, \eta, g) \equiv \text{Pr}(\bar{U} \leq V(q, \eta, g)) = \frac{V(q, \eta, g) - (2 - \beta)}{4 - (2 - \beta)} = \frac{q + \eta + \beta(1 - g) - 2}{2 + \beta} \quad (3.1)$$

[1]Note that payoffs are suppressed for ease of exposition.

Given Assumption 3 and the hiring rule in (3.1), there are eight possible cases in an informed audition (see Figure 3.1):

$$\begin{aligned}
\Pr(h|q_L,\eta_L,f) &= 0 & \Pr(h|q_L,\eta_L,m) &= \frac{\beta}{2+\beta} \\
\Pr(h|q_H,\eta_L,f) &= \frac{1}{2+\beta} & \Pr(h|q_H,\eta_L,m) &= \frac{1+\beta}{2+\beta} \\
\Pr(h|q_L,\eta_H,f) &= \frac{1}{2+\beta} & \Pr(h|q_L,\eta_H,m) &= \frac{1+\beta}{2+\beta} \\
\Pr(h|q_H,\eta_H,f) &= \frac{2}{2+\beta} & \Pr(h|q_H,\eta_H,m) &= 1
\end{aligned} \tag{3.2}$$

Importantly, the evaluator never hires a low-ability-low-effort female in an informed audition. This makes sense because, by reverting to the outside option, the evaluator can never do worse but may be lucky to engage a substitute player of higher ability, higher effort or different gender. On the other hand, the evaluator always hires a high-ability-high-effort male. This is because, by reverting to the outside option, the evaluator cannot do any better. The remaining hiring probabilities depend on the strength of the evaluator's bias. Intuitively, for females, the hiring probabilities are decreasing in β. Conversely, for males, those are increasing in β.

3.2 Effort Decision of Applicant in Stage 2

Given the above the hiring probabilities, a risk-neutral applicant chooses effort to maximise her expected payoff. Substituting for the eight possible cases in (2.2) gives the following expected utilities:

$$\begin{aligned}
U(e_L,\eta_L,f) &= -\frac{1}{2} & U(e_L,\eta_L,m) &= \frac{\beta}{2+\beta}w - \frac{1}{2} \\
U(e_H,\eta_L,f) &= \frac{1}{2+\beta}w - 2 & U(e_H,\eta_L,m) &= \frac{1+\beta}{2+\beta}w - 2 \\
U(e_L,\eta_H,f) &= \frac{1}{2+\beta}w - \frac{1}{4} & U(e_L,\eta_H,m) &= \frac{1+\beta}{2+\beta}w - \frac{1}{4} \\
U(e_H,\eta_H,f) &= \frac{2}{2+\beta}w - 1 & U(e_H,\eta_H,m) &= w - 1
\end{aligned}$$

where $w \equiv \mathbb{E}[\bar{U}] = 3 - \frac{\beta}{2}$. Because the prior that the applicant is of high ability does not enter the applicant's expected utility, I can compare participation and incentive constraints for different levels of bias to determine her effort response.

Proposition 1. *The effort responses of applicants partition the evaluator's bias $\beta \in [0,2]$ into three regions (see Figure 3.2): (i) For a low bias, $\beta \in \beta_L^I \equiv [0, \frac{5-\sqrt{17}}{2})$, only high-ability applicants participate and exert high effort. (ii) For a moderate bias, $\beta \in \beta_M^I \equiv [\frac{5-\sqrt{17}}{2}, \frac{6}{5}]$,* *low-ability males also participate and exert low effort. (iii) For a high bias, $\beta \in \beta_H^I \equiv (\frac{6}{5}, 2]$, all applicants except low-ability females participate and exert low effort.*

Proof. Low-ability females drop out of the audition for all possible biases (see Figure 3.3). High-ability females exert high effort if $\beta \in [0, \frac{6}{5}]$ and low effort if $\beta \in (\frac{6}{5}, 2]$ (see Figure 3.4). Low-ability males drop out of the audition if β is strictly lower than $\frac{5-\sqrt{17}}{2}$, and exert low effort if $\beta \in [\frac{5-\sqrt{17}}{2}, 2]$ (see Figure 3.5). High-ability males exert high effort if $\beta \in [0, \frac{6}{5}]$ and low effort if $\beta \in (\frac{6}{5}, 2]$ (see Figure 3.6). □

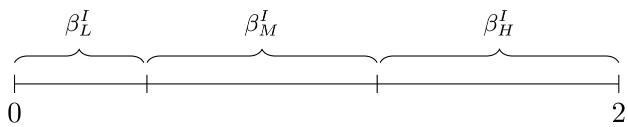

Figure 3.2: Partition of Evaluator's Bias in Informed Audition

Note the gender differences for low-ability applicants: $U(e_L, \eta_L, m)$ is increasing in the evaluator's bias while $U(e_L, \eta_L, f)$ is negative and independent of the evaluator's bias.

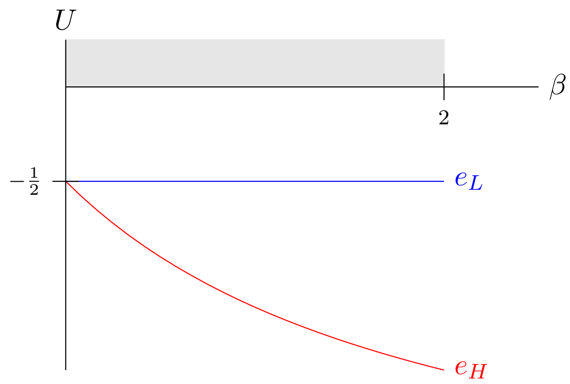

Figure 3.3: Effort Choice of Low-Ability Female in Informed Audition

3.3 Expected Utility of Evaluator in Stage 1

The evaluator's prior is that the applicant is equally likely to be a female or male when committing to the hiring process and wage w. Furthermore, he knows that the applicant is of high ability with probability $p \in (0, 1)$, and that ability is independent of gender. Hence, all four proper subgames shown in Figure 3.1 are reached with nonzero probability and are, thus, relevant when calculating the evaluator's expected gross and net utility for different values of β.

Suppose the evaluator's bias against female applicants is low; that is, $\beta \in \beta_L^I$ where β_L^I is defined as in Proposition 1(i). Then, low-ability males and females do not participate and

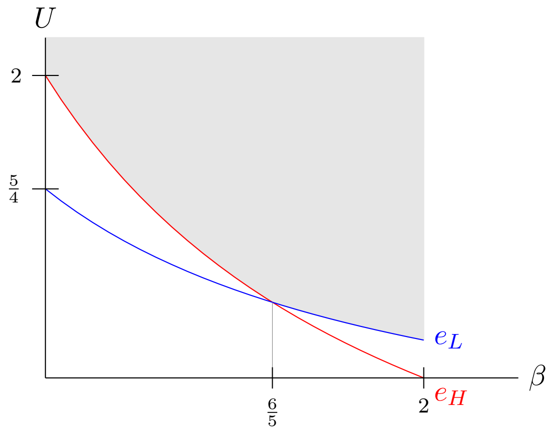

Figure 3.4: Effort Choice of High-Ability Female in Informed Audition

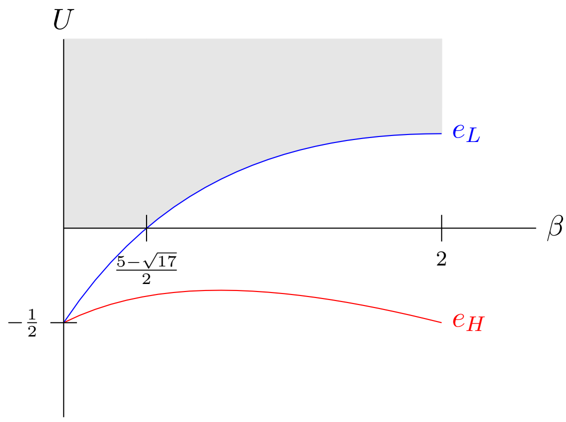

Figure 3.5: Effort Choice of Low-Ability Male in Informed Audition

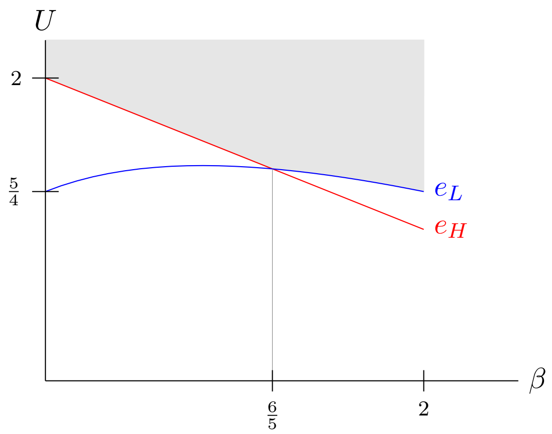

Figure 3.6: Effort Choice of High-Ability Male in Informed Audition

high-ability males and females exert high effort. In subgames (a) and (c), the evaluator, therefore, hires the outside option with probability one. With a slight abuse of notation, the evaluator's expected gross utility is the expected value of the outside option:

$$\mathbb{E}[V|O,\eta_L,f]=\mathbb{E}[V|O,\eta_L,m]=\mathbb{E}[\bar{U}]=3-\frac{\beta_L^I}{2}$$

As $w\equiv\mathbb{E}[\bar{U}]$, the evaluator, by construction, breaks even after accounting for wage costs in these two subgames. In subgame (b), the evaluator's gross utility from hiring the applicant is $4-\beta_L^I$ with corresponding hiring probability $\frac{2}{2+\beta_L^I}$. The evaluator hires the outside option with expected value $3-\frac{\beta_L^I}{2}$ with complementary probability $\frac{\beta_L^I}{2+\beta_L^I}$. The evaluator's expected gross utility in subgame (b) can then be calculated as the probability that he hires the applicant times the utility from hiring her plus the probability that he hires the outside option times the expected value of the outside option[2]:

$$\begin{aligned}\mathbb{E}[V|q_H,\eta_H,f]&=\Pr(h|q_H,\eta_H,f)V(q_H,\eta_H,f)+[1-\Pr(h|q_H,\eta_H,f)]\mathbb{E}[\bar{U}]\\&=\frac{2}{2+\beta_L^I}(4-\beta_L^I)+\frac{\beta_L^I}{2+\beta_L^I}(3-\frac{\beta_L^I}{2})\\&=2-\frac{\beta_L^I}{2}+\frac{4}{2+\beta_L^I}\end{aligned}$$

In subgame (d), the evaluator's gross utility from hiring the applicant is 4 with corresponding hiring probability 1. His expected gross utility in this subgame, therefore, is:

$$\mathbb{E}[V|q_H,\eta_H,m]=\Pr(h|q_H,\eta_H,m)V(q_H,\eta_H,m)=4$$

Under the common prior assumption, the evaluator's overall expected gross utility in stage 1 in an informed audition is:

$$\begin{aligned}\mathbb{E}[V|\beta_L^I]&=\frac{1-p}{2}\Big[\mathbb{E}[V|O,\eta_L,f]+\mathbb{E}[V|O,\eta_L,m]\Big]+\frac{p}{2}\Big[\mathbb{E}[V|q_H,\eta_H,f]+\mathbb{E}[V|q_H,\eta_H,m]\Big]\\&=\frac{1-p}{2}\Big[(3-\frac{\beta_L^I}{2})+(3-\frac{\beta_L^I}{2})\Big]+\frac{p}{2}\Big[(2-\frac{\beta_L^I}{2}+\frac{4}{2+\beta_L^I})+4\Big]\end{aligned}$$

After subtracting wage costs, the evaluator's expected net utility Π_I is:

$$\begin{aligned}\mathbb{E}[\Pi_I|\beta_L^I]&=\mathbb{E}[V|\beta_L^I]-w\\&=\frac{p}{2}\Big[\frac{2}{2+\beta_L^I}(1-\frac{\beta_L^I}{2})+(1+\frac{\beta_L^I}{2})\Big]>0\end{aligned}\tag{3.3}$$

Note that the evaluator breaks even in subgames (a) and (c), and expects positive net utility in subgames (b) and (d) if his bias against female applicants is low. Intuitively, when the evaluator is almost impartial, the outcome in an informed audition is efficient; that is,

[2] Recall that, by Assumption 2, the value of $\bar{U}$ is unknown to the evaluator when making the hiring decision. Therefore, all calculations use the *unconditional* expected value of the outside option. The model's results are robust to altering the timing and using the *conditional* expected value. However, when engaging a substitute player from an agency, it is more appropriately interpreted as a gamble as done here.

requiring a CV acts as an effective means to screen applicants by ability and to provide targeted effort incentives.

Suppose the evaluator is moderately biased against female applicants; that is, $\beta \in \beta^I_M$ where β^I_M is defined as in Proposition 1(ii). Then, subgames (a), (b) and (d) remain unchanged. Low-ability males participate and exert low effort (see Figure 3.5). The evaluator's gross utility from hiring the applicant in subgame (c) changes to 2 with corresponding hiring probability $\frac{\beta^I_M}{2+\beta^I_M}$. The evaluator hires the outside option with expected value $3 - \frac{\beta^I_M}{2}$ with complementary probability $\frac{2}{2+\beta^I_M}$. The evaluator's expected gross utility in subgame (c), therefore, becomes:

$$\begin{aligned}\mathbb{E}[V|q_L,\eta_L,m] &= \Pr(h|q_L,\eta_L,m)V(q_L,\eta_L,m) + [1-\Pr(h|q_L,\eta_L,m)]\mathbb{E}[\bar{U}] \\ &= \frac{\beta^I_M}{2+\beta^I_M}2 + \frac{2}{2+\beta^I_M}(3-\frac{\beta^I_M}{2}) \\ &= 1 + \frac{4}{2+\beta^I_M}\end{aligned}$$

which is strictly less than $\mathbb{E}[\bar{U}]$ for all β^I_M. Because $w \equiv \mathbb{E}[\bar{U}]$, the evaluator's expected utility net of wage costs is, thus, negative in subgame (c) when he is moderately biased. The evaluator's overall expected gross utility becomes:

$$\begin{aligned}\mathbb{E}[V|\beta^I_M] &= \frac{1-p}{2}\Big[\mathbb{E}[V|O,\eta_L,f] + \mathbb{E}[V|q_L,\eta_L,m]\Big] + \frac{p}{2}\Big[\mathbb{E}[V|q_H,\eta_H,f] + \mathbb{E}[V|q_H,\eta_H,m]\Big] \\ &= \frac{1-p}{2}\Big[(3-\frac{\beta^I_M}{2}) + (1+\frac{4}{2+\beta^I_M})\Big] + \frac{p}{2}\Big[(2-\frac{\beta^I_M}{2}+\frac{2}{2+\beta^I_M}) + 4\Big]\end{aligned}$$

After subtracting wage costs, the evaluator's expected net utility Π_I is:

$$\begin{aligned}\mathbb{E}[\Pi_I|\beta^I_M] &= \mathbb{E}[V|\beta^I_M] - w \\ &= \frac{1-p}{2}\left[\frac{\beta^I_M}{2+\beta^I_M}(\frac{\beta^I_M}{2}-1)\right] + \frac{p}{2}\left[\frac{2}{2+\beta^I_M}(1-\frac{\beta^I_M}{2}) + (1+\frac{\beta^I_M}{2})\right] \gtrless 0 \qquad (3.4)\end{aligned}$$

Note that the evaluator breaks even in subgame (a), and continues to expect positive net utility in subgames (b) and (d) if he is moderately biased against female applicants. In subgame (c), the evaluator instead expects negative net utility because his bias induces low-ability males to participate and exert low effort rather than dropping out. Low-ability males essentially exploit the evaluator's bias: when they know the bias to be moderate, the probability of being hired conditional on exerting low effort and identifying themselves as male is sufficiently high to make the low effort costs worthwhile.

Suppose the evaluator is highly biased against female applicants; that is, $\beta \in \beta^I_H$ where β^I_H is defined as in Proposition 1(iii). Then, subgames (a) and (c) remain unchanged. High-ability females and males exert low effort (see Figure 3.4 and Figure 3.6, respectively). The evaluator's gross utility from hiring the applicant in subgame (b) changes to $3 - \beta_H^I$ with corresponding hiring probability $\frac{1}{2+\beta_H^I}$. The evaluator hires the outside option with expected value $3 - \frac{\beta_H^I}{2}$ with complementary probability $\frac{1+\beta_H^I}{2+\beta_H^I}$. The evaluator's expected gross utility in subgame (b), therefore, falls to:

$$\begin{aligned}\mathbb{E}[V|q_L, \eta_H, f] &= \Pr(h|q_L, \eta_H, f)V(q_L, \eta_H, f) + [1 - \Pr(h|q_L, \eta_H, f)]\mathbb{E}[\bar{U}] \\ &= \frac{1}{2+\beta_H^I}(3 - \beta_H^I) + \frac{1+\beta_H^I}{2+\beta_H^I}(3 - \frac{\beta_H^I}{2}) \\ &= \frac{5}{2} - \frac{\beta_H^I}{2} + \frac{1}{2+\beta_H^I}\end{aligned}$$

The evaluator's gross utility from hiring the applicant in subgame (d) changes to 3 with corresponding hiring probability $\frac{1+\beta_H^I}{2+\beta_H^I}$. The evaluator hires the outside option with expected value $3 - \frac{\beta_H^I}{2}$ with complementary probability $\frac{1}{2+\beta_H^I}$. The evaluator's expected gross utility in subgame (d), therefore, falls to:

$$\begin{aligned}\mathbb{E}[V|q_L, \eta_H, m] &= \Pr(h|q_L, \eta_H, m)V(q_L, \eta_H, m) + [1 - \Pr(h|q_L, \eta_H, m)]\mathbb{E}[\bar{U}] \\ &= \frac{1+\beta_H^I}{2+\beta_H^I}3 + \frac{1}{2+\beta_H^I}(3 - \frac{\beta_H^I}{2}) \\ &= \frac{5}{2} + \frac{1}{2+\beta_H^I}\end{aligned}$$

The evaluator's overall expected gross utility becomes:

$$\begin{aligned}\mathbb{E}[V|\beta_H^I] &= \frac{1-p}{2}\Big[\mathbb{E}[V|O, \eta_L, f] + \mathbb{E}[V|q_L, \eta_L, m]\Big] + \frac{p}{2}\Big[\mathbb{E}[V|q_L, \eta_H, f] + \mathbb{E}[V|q_L, \eta_H, m]\Big] \\ &= \frac{1-p}{2}\Big[(3 - \frac{\beta_H^I}{2}) + (1 + \frac{4}{2+\beta_H^I})\Big] + \frac{p}{2}\Big[(\frac{5}{2} - \frac{\beta_H^I}{2} + \frac{1}{2+\beta_H^I}) + (\frac{5}{2} + \frac{1}{2+\beta_H^I})\Big]\end{aligned}$$

After subtracting wage costs, the evaluator's expected net utility Π_I is:

$$\begin{aligned}\mathbb{E}[\Pi_I|\beta_H^I] &= \mathbb{E}[V|\beta_H^I] - w \\ &= \frac{1-p}{2}\Big[\frac{\beta_H^I}{2+\beta_H^I}(\frac{\beta_H^I}{2} - 1)\Big] + \frac{p}{2}\Big[\frac{1}{2+\beta_H^I}(-\frac{\beta_H^I}{2}) + \frac{1+\beta_H^I}{2+\beta_H^I}(\frac{\beta_H^I}{2})\Big] \gtrless 0 \qquad (3.5)\end{aligned}$$

Note that the evaluator breaks even in subgame (a), and continues to expect positive net utility in subgame (d) if he is highly biased against female applicants. In subgame (b) and (c), the evaluator instead expects negative net utility. His bias is sufficiently high that high-ability females exert low rather than high effort, and low-ability males find it worthwhile to participate and exert low effort. Intuitively, when high-ability females know the evaluator's bias to be substantial, the marginal cost from putting in more effort outweighs the marginal benefit from being hired more often. They are deterred from exerting high effort at such a high level of bias. Conversely, high-ability males rest on their laurels and have little incentive to exert high effort. Their benefit from exploiting the bias and being able to identify themselves as male makes the cost increase from exerting high rather than low effort not worthwhile.

3.4 The Role of Ability

Note the role of ability in the evaluator's expected net utility. For a low bias, the evaluator's expected net utility in (3.3) is increasing in the prior that the applicant is of high ability: all other things being equal, an increase in p makes it more likely that the evaluator faces the profitable subgames (b) or (d) relative to subgames (a) or (c) in which he breaks even with the outside option. This raises the evaluator's ex-ante expected revenue from the audition. For a moderate bias, the evaluator's expected net utility in (3.4) is also increasing in p. However, the benefit of an increasingly skilled applicant pool is more pronounced: as subgame (c) becomes less likely, the evaluator is able to avoid the losses from potentially hiring low-ability males who are attracted to participate in the audition. For a high bias, the increase in the evaluator's expected net utility in (3.5) is more attenuated: as p increases, the evaluator faces both the unprofitable subgame (b) and the profitable subgame (d) more often.

4

Blind Audition

In a blind audition, the evaluator does not learn the applicant's type $\theta = (\eta, g)$ in stage 1. He only observes the applicant's performance quality q in stage 3. Thus, there are no proper subgames but two information sets: q_L and q_H. Each information set contains four nodes when the evaluator has to make the hiring decision in $s = 3$ (see Figure 4.1). Because these information sets are non-singletons, stage 2 and 3 in a blind audition can be viewed as a game of incomplete information. Therefore, I need to specify beliefs for the evaluator; that is, a probability distribution over the nodes in both q_L and q_H. Intuitively, this probability distribution constitutes the evaluator's revised beliefs how likely it is upon hearing a low- or high-quality performance, that he is facing a particular type of applicant, such as a high-ability female, behind the curtain. In what follows, I focus on the case that the bias is either sufficiently low to induce all applicants to exert high effort or sufficiently high to induce only high-ability applicants to exert high effort[1]. I refer to the former case as pooling and the latter case as separating.

4.1 Pooling Equilibrium

Suppose the evaluator's bias is sufficiently low to induce all applicants to exert high effort. However, when observing a low performance quality off the equilibrium path, the evaluator believes the applicant to be of low ability. Upon observing a low-quality performance, the evaluator can, therefore, infer the applicant's ability from her action and expects gross

[1]Note that I do not consider the case in which applicants pool on low effort because the focus of the thesis is on whether, and under which conditions, blind auditions *can* be superior. To see why this can never obtain in a low-effort pooling equilibrium, compare it to the worst-case scenario in an informed audition; that is, for $\beta \in \beta_H^I$, all applicants except low-ability females exert low effort. This scenario gives the evaluator still a higher payoff because attracting low-ability-low-effort females is never profitable. The evaluator could have broken even with the outside option and, therefore, is worse off.

[2]Note that payoffs are suppressed for ease of exposition.

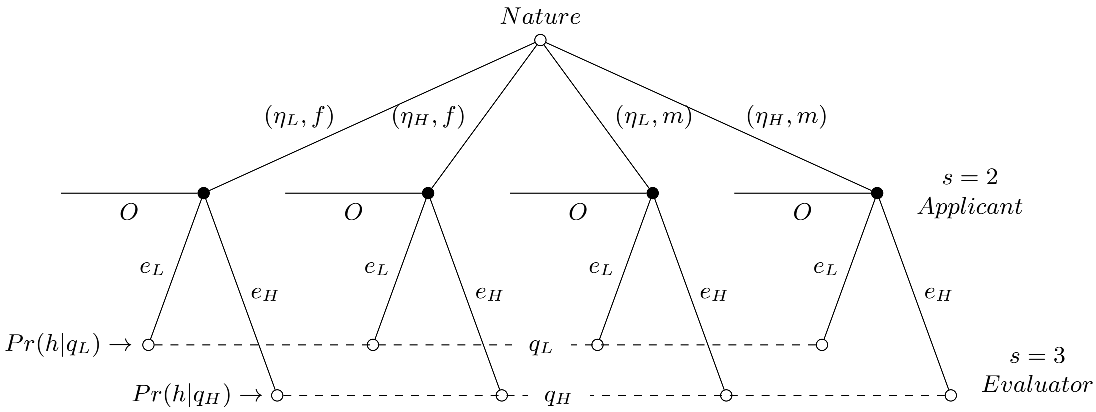

Figure 4.1: Two Information Sets in Blind Audition[2]

utility[3]:

$$\mathbb{E}[V|q_L] = q_L + 1 - \frac{\beta}{2} = 2 - \frac{\beta}{2}$$

from hiring the applicant. Upon observing a high-quality performance, the evaluator cannot infer the applicant's ability from her action and holds a belief that is equal to his prior; that is, $\Pr(\eta_H) = p \in (0, 1)$. He expects gross utility:

$$\mathbb{E}[V|q_H] = q_H + \mathbb{E}[\eta|q_H] - \frac{\beta}{2} = 2 + [(1-p) + 2p] - \frac{\beta}{2} = 3 + p - \frac{\beta}{2}$$

from hiring the applicant. The evaluator can revert to the random outside option $\bar{U} \sim U[2-\beta, 4]$. He, therefore, hires the applicant at each information set with a probability derived from comparing the expected utility from hiring an applicant to the cumulative distribution function of $\bar{U}$:

$$\Pr(\text{hired}|q) \equiv \Pr(\bar{U} \leq \mathbb{E}[V|q]) = \frac{\mathbb{E}[V|q] - (2-\beta)}{4 - (2-\beta)} = \frac{q + \mathbb{E}[\eta|q] - \frac{\beta}{2} - (2-\beta)}{2+\beta} \tag{4.1}$$

Given Assumption 3 and the hiring rule in (4.1), there are two possible cases for both female and male applicants given by:

$$\Pr(h|q_L) = \Pr(\bar{U} \leq 2 - \frac{\beta}{2}) = \frac{\beta}{4+2\beta}$$
$$\Pr(h|q_H) = \Pr(\bar{U} \leq 3 + p - \frac{\beta}{2}) = \frac{2(1+p)+\beta}{4+2\beta} \tag{4.2}$$

Compared to an informed audition, female and male applicants now face the same hiring probabilities as the evaluator cannot bias gender. Similarly, high- and low-ability applicants face the same hiring probabilities as they cannot identify themselves in a pooling blind audition.

[3]Note that the evaluator cannot infer gender in a pooling or separating blind audition because, with η fixed, females and males have the same payoff structure. In effect, the evaluator now faces only two types of applicants, η_L and η_H.

4.1.1 Effort Decision of Applicant in Stage 2

Substituting for the two possible cases in the applicant's expected payoff in (2.2), and noting that ability η still plays a role via the applicant's effort costs, gives the following expected utilities for both female and male applicants in a pooling blind audition:

$$\begin{aligned} U(e_L, \eta_L) &= \frac{\beta}{4+2\beta}w - \frac{1}{2} \\ U(e_H, \eta_L) &= \frac{2(1+p)+\beta}{4+2\beta}w - 2 \\ U(e_L, \eta_H) &= \frac{\beta}{4+2\beta}w - \frac{1}{4} \\ U(e_H, \eta_H) &= \frac{2(1+p)+\beta}{4+2\beta}w - 1 \end{aligned}$$

where $w \equiv \mathbb{E}[\bar{U}] = 3 - \frac{\beta}{2}$. Note that, as the prior that the applicant is of high ability increases, the expected gross utility from hiring an applicant at q_H increases. This is because, ceteris paribus, an increasingly skilled applicant pool raises the evaluator's expected revenue from hiring at the information set. Hence, the evaluator finds it optimal to hire more often at q_H. Graphically, an increase in p implies an upward shift of the high-effort curve for both low- and high-ability applicants in (β, U)-space (see Figure 4.2 and 4.3). I now show that p determines whether a high-effort pooling equilibrium can be supported, and that p places an upper bound on the evaluator's bias in any such pooling equilibrium.

Proposition 2. *(i) For $p < \frac{1}{3}$, no high-effort pooling equilibrium exists. (ii) For $p = \frac{1}{3}$, a high-effort pooling equilibrium exists if the evaluator is unbiased. (iii) For $\frac{1}{3} < p < 1$, a high-effort pooling equilibrium exists if $\beta \in \beta_L^B \equiv [0, \beta(p)]$, where $\beta'(p) > 0$.*

Proof. (i) For $p < \frac{1}{3}$, I show that $U(e_H, \eta_L)$ is strictly negative for all possible values of β. It is always better for low-ability applicants to not participate. First, note that $U(e_H, \eta_L)$ is strictly decreasing in β:

$$\frac{\partial U(e_H, \eta_L)}{\partial \beta} = -\frac{4p}{(4+2\beta)^2}(3 - \frac{\beta}{2}) - \frac{2(1+p)+\beta}{8+4\beta} < 0 \; \forall \beta \in [0, 2]$$

This implies that if $U(e_H, \eta_L) < 0$ at $\beta = 0$, then $U(e_H, \eta_L) < 0$ for all $\beta \in [0, 2]$. I can use this condition to solve for the corresponding prior:

$$U(e_H, \eta_L)\Big|_{\beta=0} = \frac{1+p}{2}3 - 2 < 0 \Rightarrow p < \frac{1}{3}$$

Thus, for $p < \frac{1}{3}$, pooling on high effort is never possible.

(ii) For $p = \frac{1}{3}$, $U(e_H, \eta_L) = 0$ at $\beta = 0$. By the previous argument, $U(e_H, \eta_L) < 0$ for all $\beta \in (0, 2]$. This implies that low-ability applicants are indifferent between exerting high effort and not participating only if the evaluator is unbiased. Furthermore, it is optimal for high-ability applicants to exert high effort whenever it is optimal for low-ability applicants to exert high effort. This is because $U(e_L, \cdot)$ shifts up by $\frac{1}{4}$ while $U(e_H, \cdot)$ shifts up by 1 in the move from η_L to η_H (see Figure 4.2 and 4.3). Thus, pooling on high effort is possible at $p = \frac{1}{3}$ if $\beta = 0$.

(iii) For $\frac{1}{3} < p < 1$, there exists a range of biases for which $U(e_H, \eta_L)$ is positive. The upper bound on β for which a pooling equilibrium can be supported is determined by the indifference condition:

$$U(e_H, \eta_L) = \frac{2(1+p)+\beta}{4+2}(3 - \frac{\beta}{2}) - 2 = 0$$

Given that $\beta \in [0, 2]$, the above equation is uniquely solved by $\beta(p) \equiv \sqrt{p^2 + 16p} - p - 2$. The upper bound $\beta(p)$ is strictly increasing in p:

$$\frac{d\beta(p)}{dp} = \frac{p+8}{\sqrt{p(p+16)}} - 1 > 0 \; \forall \; 0 < p < 1$$

□

Lemma 1. *For $\frac{1}{3} < p < 1$, as $p \uparrow 1$, the upper bound on the pooling equilibrium $\beta(p)$ is approaching the infimum of the separating equilibrium (see Proposition 3).*

Proof. Taking the limit of $\beta(p)$ yields:

$$\lim_{p \uparrow 1} \beta(p) = \sqrt{17} - 3 > 0$$

□

With the observation that it is always optimal for high-ability applicants to exert high effort if it is optimal for low-ability applicants to exert high effort, pooling on high effort is possible if the prior that the applicant is of high ability is sufficiently high and the evaluator's bias is not too large. Furthermore, an increasingly able applicant pool makes it easier for for the fewer low-ability to hide behind the more-and-more high-ability applicants.

4.1.2 Expected Utility of Evaluator in Stage 1

While the evaluator does not learn the applicant's type, the applicant's payoff structure is common knowledge. In particular, for the evaluator's belief that all applicants exert high effort to be consistent, he needs to have a sufficiently low bias; that is, $\beta \in \beta_L^B$ with $\frac{1}{3} \leq p < 1$. His expected utility from hiring the applicant at the information set q_H is $3 + p - \frac{\beta_L^B}{2}$ with corresponding hiring probability $\frac{2(1+p)+\beta_L^B}{4+2\beta_L^B}$. The evaluator hires the outside option with expected value $3 - \frac{\beta_L^B}{2}$ with complementary probability $\frac{2(1-p)+\beta_L^B}{4+2\beta_L^B}$. Given the

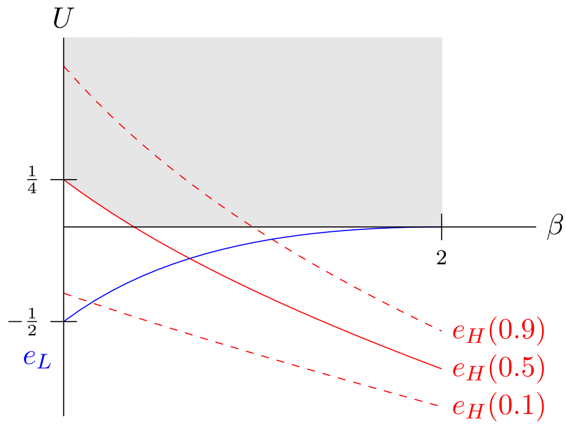

Figure 4.2: Effort Choice of Low-Ability Applicant in Pooling Blind Audition for Different Priors

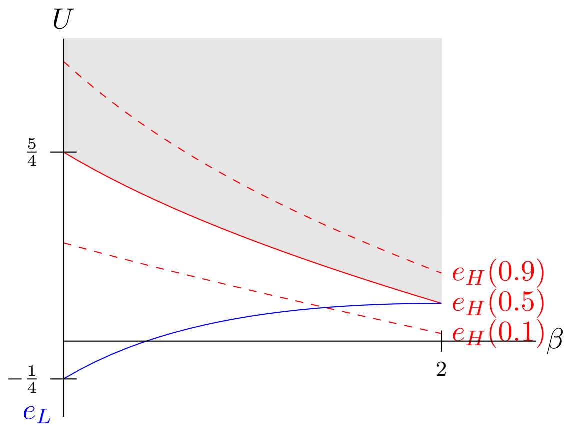

Figure 4.3: Effort Choice of High-Ability Applicant in Pooling Blind Audition

evaluator's beliefs, his overall expected gross utility in stage 1 in a pooling blind audition is:

$$\begin{aligned}\mathbb{E}[V|\beta_L^B] &= \Pr(h|q_H)\mathbb{E}[V|q_H] + [1 - \Pr(h|q_H)]\mathbb{E}[\bar{U}] \\ &= \frac{2(1+p)+\beta_L^B}{4+2\beta_L^B}(3+p-\frac{\beta_L^B}{2}) + \frac{2(1-p)+\beta_L^B}{4+2\beta_L^B}(3-\frac{\beta_L^B}{2})\end{aligned}$$

After subtracting wage costs, the evaluator's expected net utility Π_B is:

$$\begin{aligned}\mathbb{E}[\Pi_B|\beta_L^B] &= \mathbb{E}[V|\beta_L^B] - w \\ &= \frac{2(1+p)+\beta_L^B}{4+2\beta_L^B}\left[3+p-\frac{\beta_L^B}{2} - (3-\frac{\beta_L^B}{2})\right] \\ &= \frac{2(1+p)+\beta_L^B}{4+2\beta_L^B}p > 0 \qquad (4.3)\end{aligned}$$

Note that (4.3) is increasing in the prior that the applicant is of high ability because, ceteris paribus, an increasingly skilled applicant pool raises the evaluator's expected revenue from hiring in a pooling blind audition. Intuitively, the revised belief of the evaluator upon observing a high-quality performance takes into account that it is ex-ante more likely to face a high-ability female or a high-ability male behind the curtain; that is, the evaluator is more likely to be at the second or fourth node from the left of the information set q_H in Figure 4.1.

4.2 Separating Equilibrium

Suppose the evaluator's bias is sufficiently high that only high-ability applicants exert high effort. As before, the evaluator believes the applicant to be of low ability when observing a low performance quality. Therefore, the evaluator has degenerate posterior beliefs at both information sets. Upon observing a low-quality performance, he holds the belief that the applicant is of low ability with probability one, and male or female with equal probability. Upon observing a high-quality performance, he believes the applicant to be of high ability with probability one, and male or female with equal probability. At q_L, the evaluator expects gross utility:

$$\mathbb{E}[V|q_L] = q_L + 1 - \frac{\beta}{2} = 2 - \frac{\beta}{2}$$

from hiring the applicant. At q_H, the evaluator can now infer the applicant's ability from her action and expects gross utility:

$$\mathbb{E}[V|q_H] = q_H + 2 - \frac{\beta}{2} = 4 - \frac{\beta}{2}$$

from hiring the applicant. The evaluator hires the applicant at each information set with a probability derived from comparing $\mathbb{E}[V|q]$ to the cumulative distribution function of $\bar{U}$. There are two possible cases for both female and male applicants given by:

$$\begin{aligned} \Pr(h|q_L) &= \Pr(\bar{U} \leq 2 - \frac{\beta}{2}) = \frac{\beta}{4 + 2\beta} \\ \Pr(h|q_H) &= \Pr(\bar{U} \leq 4 - \frac{\beta}{2}) = \frac{4 + \beta}{4 + 2\beta} \end{aligned} \tag{4.4}$$

4.2.1 Effort Decision of Applicant in Stage 2

Substituting for the two possible cases in equation (2.2) gives the following expected utilities for both male and female applicants in a separating blind audition:

$$U(e_L,\eta_L) = \frac{\beta}{4+2\beta}w - \frac{1}{2}$$
$$U(e_H,\eta_L) = \frac{4+\beta}{4+2\beta}w - 2$$
$$U(e_L,\eta_H) = \frac{\beta}{4+2\beta}w - \frac{1}{4}$$
$$U(e_H,\eta_H) = \frac{4+\beta}{4+2\beta}w - 1$$

where $w \equiv \mathbb{E}[\bar{U}] = 3 - \frac{\beta}{2}$. Because the prior that the applicant is of high ability does not enter the applicant's expected utility, I can compare participation and incentive constraints as in an informed audition to characterise equilibrium in terms of the bias parameter.

Proposition 3. *For $\beta \in \beta_H^B \equiv (\sqrt{17}-3, 2)$, a separating equilibrium exists in which only high-ability applicants exert high effort and low-ability applicants do not participate.*

Proof. For $\beta \in [0, \sqrt{17}-3]$, low-ability applicants exert high effort (see Figure 4.4). For $\beta \in (\sqrt{17}-3, 2)$, low-ability applicants do not participate in the audition (see Figure 4.4). For $\beta \in [0, 2]$, high-ability applicants exert high effort in the audition (see Figure 4.5). □

Intuitively, separation is possible if the bias is sufficiently high so that it is too costly for low-ability applicants to participate. Furthermore, note that, for $\beta < 2$, the outside option dominates low effort if the applicant is of low ability. For $\beta \leq 2$, high effort dominates low effort if the applicant is of high ability.

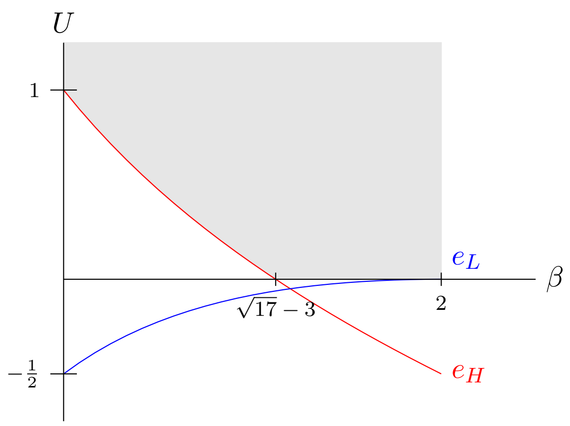

Figure 4.4: Effort Choice of Low-Ability Applicant in Separating Blind Audition

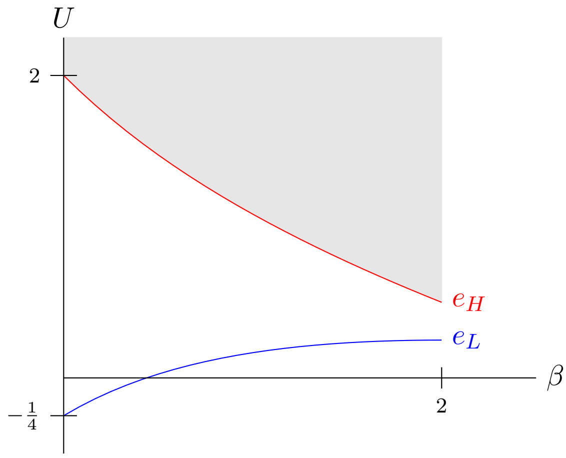

Figure 4.5: Effort Choice of High-Ability Applicant in Separating Blind Audition

4.2.2 Expected Utility of Evaluator in Stage 1

For the evaluator's belief that only high-ability applicants exert high effort to be consistent, he needs to have a sufficiently high bias β_H^B. In this case, low-ability applicants do not participate and q_L is not be reached. Therefore, the evaluator's belief at this information set is not determined by the applicant's equilibrium play. I provide a robustness check in Appendix B: two refinements based on either deviation payoffs or the D1-Criterion show that the evaluator's belief that the applicant is of low ability when observing a deviation to q_L constitutes a reasonable restriction.

The evaluator's expected utility from hiring the applicant at the information set q_H is $4 - \frac{\beta_H^B}{2}$ with corresponding hiring probability $\frac{4+\beta_H^B}{4+2\beta_H^B}$. The evaluator hires the outside option with expected value $3 - \frac{\beta_H^B}{2}$ with complementary probability $\frac{\beta_H^B}{2+\beta_H^B}$. If the applicant is a low-ability male or female, the evaluator hires the outside option with probability one as those types do not participate. In this case, the evaluator's expected utility is the expected value of the outside option:

$$\mathbb{E}[V|O] = \mathbb{E}[\bar{U}] = 3 - \frac{\beta_H^B}{2}$$

As $w \equiv \mathbb{E}[\bar{U}]$, the evaluator breaks even by construction after accounting for wage costs. Given the evaluator's beliefs, his overall expected gross utility in stage 1 in a separating blind audition is:

$$\begin{aligned}\mathbb{E}[V|\beta_H^B] &= p\Big[\Pr(h|q_H)\mathbb{E}[V|q_H] + [1 - \Pr(h|q_H)]\mathbb{E}[\bar{U}]\Big] + (1-p)\Big[\mathbb{E}[\bar{U}]\Big] \\ &= p\left[\frac{4+\beta_H^B}{4+2\beta_H^B}(4 - \frac{\beta_H^B}{2}) + \frac{\beta_H^B}{2+\beta_H^B}(3 - \frac{\beta_H^B}{2})\right] + (1-p)\left[3 - \frac{\beta_H^B}{2}\right]\end{aligned}$$

After subtracting wage costs, the evaluator's expected net utility Π_B is:

$$\begin{aligned}\mathbb{E}[\Pi_B|\beta_H^B] &= \mathbb{E}[V|\beta_H^B] - w \\ &= p\Bigg[\frac{4+\beta_H^B}{4+2\beta_H^B}\Bigg(4 - \frac{\beta_H^B}{2} - (3 - \frac{\beta_H^B}{2})\Bigg)\Bigg] \\ &= p\frac{4+\beta_H^B}{4+2\beta_H^B} > 0 \end{aligned} \tag{4.5}$$

As before, (4.5) is increasing in the prior that the applicant is of high ability. However, the channel through which an increase in the prior leads to an increase in the evaluator's profit is less pronounced. This is because, in a pooling blind audition, an increase in the prior necessarily also implied a reduction in the probability that, upon observing a high-quality performance, the evaluator faced a low-ability female or low-ability male behind the curtain; that is, the evaluator was less likely to be at the first or third node from the left of the information set q_H in Figure 4.1.

5

Comparison of Blind and Informed Audition

Having discussed the evaluator's net utility in both forms of audition, I can conclude under which conditions on the bias parameter β and the prior p the evaluator prefers a blind or informed audition to maximise his expected net utility.

Proposition 4. *(i) For $\frac{1}{3} \leq p < 1$, if the evaluator's bias against female applicants is low, $\beta \in \beta_L^B$, he prefers an informed audition over a pooling blind audition. (ii) For $0 < p < 1$, if the evaluator's bias against female applicants is high, $\beta \in \beta_H^I \cap \beta_H^B$, he prefers a separating blind audition. (iii) For $0 < p < \frac{2-\beta}{2}$, if the evaluator's bias against female applicants is moderate, $\beta \in \beta_M^I \cap \beta_H^B$, he prefers a separating blind audition. For $\frac{2-\beta}{2} < p < 1$, if the evaluator's bias against female applicants is moderate, $\beta \in \beta_M^I \cap \beta_H^B$, he prefers an informed audition.*

Proof. (i) For $\beta \in \beta_L^B \equiv [0, \sqrt{p^2 + 16p} - p - 2]$, I show that the evaluator's expected net utility in a blind audition is strictly lower than his expected net utility in an informed audition. First, from Lemma 1, it follows that $\frac{5-\sqrt{17}}{2} < \lim_{p \uparrow 1} \beta(p) < \frac{5}{4}$; that is, the largest bias for which a pooling equilibrium can be supported lies in the interval corresponding to β_M^I in an informed audition (see Figure 5.1a for an example). Therefore, I need to compare $\mathbb{E}[\Pi_B|\beta_L^B]$ with $\mathbb{E}[\Pi_I|\beta_L^I]$ and $\mathbb{E}[\Pi_I|\beta_M^I]$ to make the above conclusion:

$$\mathbb{E}[\Pi_I|\beta_L^I] > \mathbb{E}[\Pi_B|\beta_L^B]$$
$$\Rightarrow \frac{p}{2}\left[\frac{2}{2+\beta_L^I}(1 - \frac{\beta_L^I}{2}) + (1 + \frac{\beta_L^I}{2})\right] > \frac{2(1+p) + \beta_L^B}{4 + 2\beta_L^B}p$$

which is true for any $0 < p < 1$ and $\beta > -2$.

$$\mathbb{E}[\Pi_I|\beta_M^I] > \mathbb{E}[\Pi_B|\beta_L^B]$$
$$\Rightarrow \frac{1-p}{2}\left[\frac{\beta_M^I}{2+\beta_M^I}(\frac{\beta_M^I}{2} - 1)\right] + \frac{p}{2}\left[\frac{2}{2+\beta_M^I}(1 - \frac{\beta_M^I}{2}) + (1 + \frac{\beta_M^I}{2})\right] > \frac{2(1+p) + \beta_L^B}{4 + 2\beta_L^B}p$$

which is true for any $\frac{1}{5} < p < 1$ and $\beta > -2$. By Proposition 2, the lower bound on the prior is more restrictive to support any pooling equilibrium. By assumption, the evaluator's bias is non-negative. Therefore, both inequalities are always satisfied and the evaluator's expected net utility in a blind audition is strictly lower.

(ii) For $\beta \in \beta_H^I \cap \beta_H^B = (\frac{6}{5}, 2)$, I show that the evaluator's net utility in a blind audition is strictly greater than in an informed audition:

$$\mathbb{E}[\Pi_B|\beta_H^B] > \mathbb{E}[\Pi_I|\beta_H^I]$$
$$\Rightarrow p\frac{4+\beta_H^B}{4+2\beta_H^B} > \frac{1-p}{2}\left[\frac{\beta_H^I}{2+\beta_H^I}(\frac{\beta_H^I}{2}-1)\right] + \frac{p}{2}\left[\frac{1}{2+\beta_H^I}(-\frac{\beta_H^I}{2}) + \frac{1+\beta_H^I}{2+\beta_H^I}(\frac{\beta_H^I}{2})\right]$$

which is true for any $-\frac{1}{8} < p \leq 1$ and $1-\sqrt{8p+1} < \beta < \sqrt{8p+1}+1$. By assumption, the lower and upper bound on the prior are more restrictive, and the evaluator's bias lies in the interval $[0, 2]$. Therefore, the above inequality is always satisfied.

(iii) For $\beta \in \beta_M^I \cap \beta_H^B = (\sqrt{17}-3, \frac{6}{5}]$, I show that the evaluator's net utility in a blind audition is strictly greater than his expected utility in an informed audition only if $0 < p < \frac{2-\beta}{2}$:

$$\mathbb{E}[\Pi_B|\beta_H^B] > \mathbb{E}[\Pi_I|\beta_M^I]$$
$$\Rightarrow p\frac{4+\beta_H^B}{4+2\beta_H^B} > \frac{1-p}{2}\left[\frac{\beta_M^I}{2+\beta_M^I}(\frac{\beta_M^I}{2}-1)\right] + \frac{p}{2}\left[\frac{2}{2+\beta_M^I}(1-\frac{\beta_M^I}{2}) + (1+\frac{\beta_M^I}{2})\right]$$

which is true for any $p \leq 1$ and $0 < \beta < 2-2p$. Conversely, the strict reverse inequality holds for any $p \leq 1$ and $\beta > 2-2p$. □

Intuitively, the evaluator is better off having more information if his bias against female applicants is low. First, an informed audition allows him to attract a pool of exclusively high-ability applicants. Moving to a blind audition would also induce low-ability applicants to participate and exert high effort if the evaluator's prior is sufficiently high (see Figure 5.1a and 5.1b for an example). Second, a blind audition would also imply that the evaluator can no longer hire a high-ability male with probability one when exerting high effort as this type of applicant cannot distinguish himself from the others. However, reverting to the outside option with positive probability when high-ability males exert high effort can only make the evaluator worse off due to the upper bound on $\bar{U}$. This prediction of my model is in line with Holland (1981) observing that "equal opportunity has unloosed such an avalanche of auditionees that standards of selection and fairness are sometimes actually lowered" and that "[a]s a result, top talent is sometimes lost in the shuffle".

On the other hand, the evaluator is better off having no information about the applicant's type if he is highly biased against female applicants. In this case, a blind audition provides targeted effort incentives. In an informed audition, in contrast, the evaluator's bias would

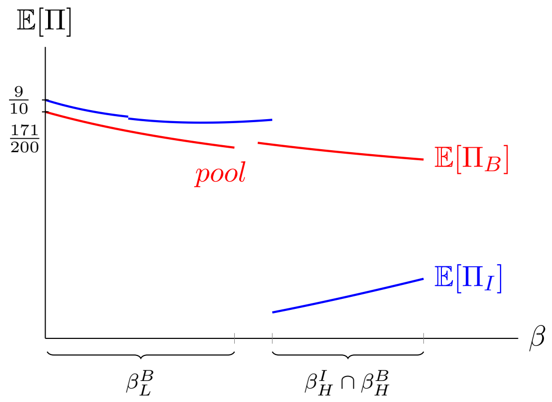

(a) $\Pr(\eta_H) = 0.9$

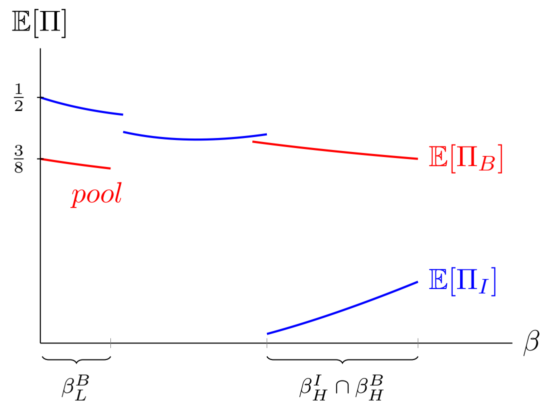

(b) $\Pr(\eta_H) = 0.5$

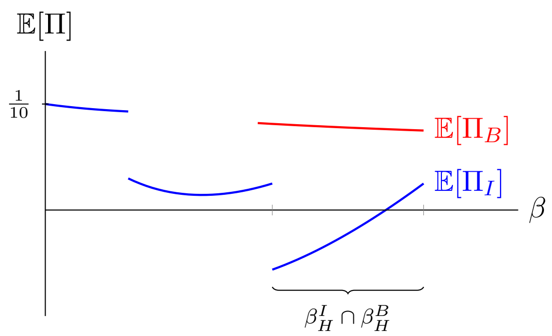

(c) $\Pr(\eta_H) = 0.1$

Figure 5.1: Evaluator's Expected Net Utility in Blind and Informed Audition

distort effort incentives: all applicants except low-ability females would exert low effort. This prediction of my model echoes the insight of Gladwell (2005) regarding first impression biases; that is, "by fixing the first impression at the heart of the audition - by judging purely on the basis of ability - orchestras now hire better musicians, and better musicians mean better music" (p.253).

For a small range of moderate biases, $\beta \in \beta_M^I \cap \beta_H^B$, both a blind and informed audition

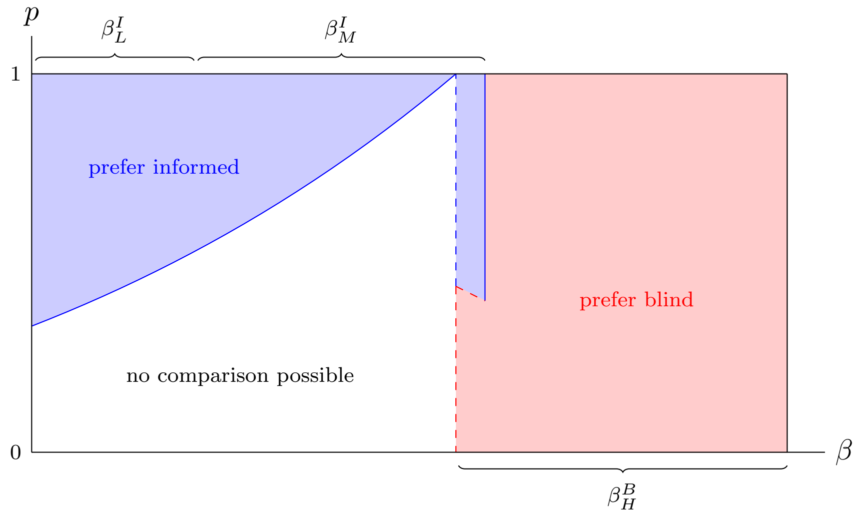

Figure 5.2: Evaluator's Preferences over Auditions for Different Biases and Priors

provide high effort incentives for high-ability applicants. However, the evaluator faces the distortion that an informed audition also attracts low-ability males, who exert only low effort. The prior determines the size of the loss from potentially hiring those types, and, thus, whether the evaluator may be better off in a separating blind audition. First, for a range of priors from $\frac{2}{5}$ to $\frac{5-\sqrt{17}}{2}$, the evaluator's preferences over auditions are non-monotonic in his bias (see Figure 5.2). Intuitively, this is because (3.4) is upward-sloping whereas (4.5) is downward-sloping on the interval $\beta_M^I \cap \beta_H^B$. However, if the pool of applicants is of predominantly low ability, the evaluator always prefers a blind audition for this range of moderate biases: it is an effective means to provide targeted incentives while reaping the benefits from breaking even with the outside option (see Figure 5.1c for an example). The evaluator avoids substantial losses which he would have incurred in an informed audition from hiring low-ability males relatively frequently. As the pool of applicants becomes increasingly able, however, this concern becomes less important. In particular, for a sufficiently high prior, the benefit from being able to hire high-ability females and high-ability males with differing probabilities outweighs the loss from potentially hiring a low-ability-low-effort male in an informed audition (see Figure 5.1a and 5.1b for an example).

To gain some intuition behind the evaluator's distribution of preferences, suppose that the bias and the prior on high ability are drawn uniformly from their respective supports. In other words, all points in Figure 5.2 are equally likely before the three-stage game is played. Therefore, I can use integration to calculate the ex-ante percentage with which a blind audition (red area) is preferred relative to an informed audition (blue area). In so doing, I discount the white areas for which no comparison is possible. The evaluator prefers a blind audition approximately sixty-four percent and an informed audition only approximately thirty-six percent of the time. At a macroeconomic level, this simple form of heterogeneity in the bias across evaluators and in the ability composition across applicant pools rationalises the observed co-existence of blind and informed auditions in the economy. Furthermore, it rationalises that the majority of U.S. symphony orchestras use a blind audition as an outcome of optimisation (Goldin and Rouse, 2000).

5.1 Applicant Preferences over Auditions

To complement the above discussion of the evaluator's optimal audition form, I consider the applicant's preferences over auditions. In particular, having solved for optimal effort levels in both forms of audition, I can conclude under which conditions on the bias β and prior p each type prefers a blind or an informed audition to maximise her expected utility.

Proposition 5. *(i) A low-ability female prefers a blind audition if the evaluator's bias against female applicants is low; that is, $\beta \in \beta_L^B$. (ii) She is indifferent if the evaluator's bias against female applicants is high; that is, $\beta \in \beta_H^B$.*

Proof. (i) For $\beta \in \beta_L^B \equiv [0, \sqrt{p^2 + 16p} - p - 2]$ and $\frac{1}{3} \leq p < 1$, the applicant's expected utility is weakly greater in an blind audition:

$$U(e_H, \eta_L) \geq U(O, \eta_L, f)$$
$$\Rightarrow \frac{2(1+p)+\beta}{4+2\beta} w - 2 \geq 0$$

By Proposition 2, the above inequality is the constraint on the range of β for which a pooling equilibrium can be supported and, therefore, always satisfied.

(ii) For $\beta \in \beta_H^B \equiv (\sqrt{17} - 3, 2)$, the applicant does not participate in either of the auditions and her expected utility is zero, as shown in Figure 5.3. □

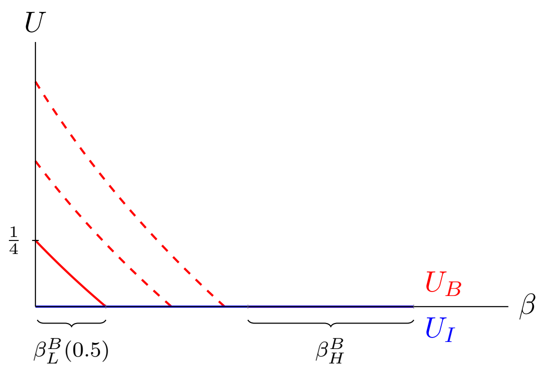

Figure 5.3: Expected Utility of Low-ability Female in Blind and Informed Audition

Intuitively, a low-ability female benefits from the move to a blind audition under the condition that the applicant pool is predominantly of high ability and the evaluator's bias is sufficiently low. It allows her to pool with high-ability applicants, concealing her low ability and gender, and, thus, be hired with nonzero probability. In an informed audition, she would never participate because having to reveal both gender and ability would play out against her.

Proposition 6. *(i) A high-ability female prefers an informed audition if the evaluator's bias against female applicants is low. (ii) She prefers a blind audition if the evaluator's bias against female applicants is high; that is,* $\beta \in \beta_H^B$.

Proof. (i) First, note that, for $\beta \in [0, 2(1-p)]$, the applicant's expected utility is weakly greater in an informed audition:

$$U(e_H, \eta_H, f) \geq U(e_H, \eta_H)$$
$$\Rightarrow \frac{2}{2+\beta} \geq \frac{2(1+p)+\beta}{4+2\beta}$$
$$\Rightarrow \beta \leq 2(1-p)$$

However, because the upper bound $\beta(p) \equiv \sqrt{p^2+16p}-p-2$, for which a pooling equilibrium can be supported, may be more restrictive than the above inequality, I need to find the range of priors for which either constraint is more restrictive. The indifference condition is:

$$\sqrt{p^2+16p}-p-2 = 2(1-p)$$
$$\Rightarrow p^* = \frac{25-\sqrt{561}}{2} \approx 0.65728$$

For priors below p^*, $\beta(p)$ is the more restrictive constraint. For priors above p^*, the upper bound for which a high-ability female prefers an informed audition is given by $\beta \leq 2(1-p)$.

(ii) For $\beta \in \beta_H^B \equiv (\sqrt{17}-3, 2)$, the applicant's expected utility is strictly greater in a blind audition, as shown in Figure 5.4. □

To gain some insight, fix a prior that is sufficiently high such that the range of biases for which a high-ability female prefers an informed audition is determined by the intersection of U_B and U_I (see Figure 5.4). Then, as the prior that the applicant is of high ability approaches one, the range of biases for which a high-ability female prefers an informed audition is shrinking and approaches zero. Intuitively, for this applicant, it becomes relatively more important to conceal her gender as the evaluator already knows that she is likely to be of high ability. Submitting a CV prior to the audition is of little value to her.

Proposition 7. *(i) A low-ability male prefers a blind audition if the evaluator's bias against female applicants is low. (ii) He prefers an informed audition if the evaluator's bias against female applicants is high; that is,* $\beta \in \beta_H^B$.

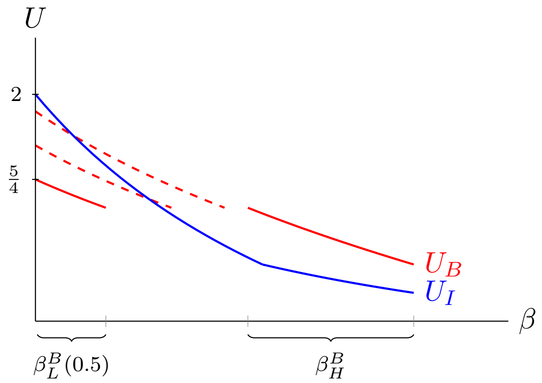

Figure 5.4: Expected Utility of High-ability Female in Blind and Informed Audition

Proof. (i) First, note that, for $\beta \in [0, \frac{5-\sqrt{17}}{2})$, the applicant does not participate in an informed audition. If $\beta \geq \frac{5-\sqrt{17}}{2}$, the applicant exerts low effort in an informed audition. Therefore, I need to compare:

$$U(e_H, \eta_L) \geq U(e_L, \eta_L, m)$$
$$\Rightarrow \frac{2(1+p)+\beta}{4+2\beta}w - 2 \geq \frac{\beta}{2+\beta}w - \frac{1}{2}$$
$$\Rightarrow \beta \leq -\sqrt{p^2+2p+49} + p + 7$$

However, because the upper bound $\beta(p) \equiv \sqrt{p^2+16p} - p - 2$ for which a pooling equilibrium can be supported may be more restrictive than the above inequality, I need to find the range of priors for which either constraint is more restrictive. The indifference condition is:

$$\sqrt{p^2+16p} - p - 2 = -\sqrt{p^2+2p+49} + p + 7$$
$$\Rightarrow p^{**} = \frac{31 - 7\sqrt{17}}{4} \approx 0.53457$$

For priors below p^{**}, $\beta(p)$ is the more restrictive constraint. For priors above p^{**}, the upper bound for which a low-ability male prefers a blind audition is given by $\beta \leq -\sqrt{p^2+2p+49} + p + 7$.

(ii) For $\beta \in \beta_H^B \equiv (\sqrt{17} - 3, 2)$, the applicant's expected utility is strictly greater in an informed audition, as shown in Figure 5.5. □

For a very low bias, the intuition behind Proposition 7 is akin to that for low-ability females. A low-ability male benefits from the move to a blind audition under the condition that the applicant pool is predominantly of high ability and the evaluator's bias is sufficiently low. For this applicant, it is attractive to conceal his low ability in a pooling blind audition. He is, in essence, benefiting from an informational externality because the pool of applicants is increasingly able. However, beyond a critical value of the bias, given by $\beta(p)$ for priors

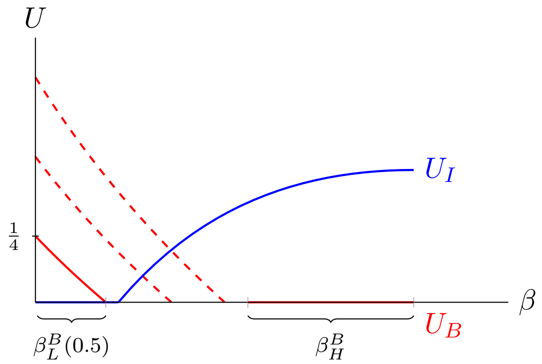

Figure 5.5: Expected Utility of Low-ability Male in Blind and Informed Audition

below and by $\beta \leq -\sqrt{p^2+2p+49}+p+7$ for priors above p^{**}, he trades off this advantage and rather identifies himself as low-ability male to reap the benefits of the evaluator's bias against female applicants.

Proposition 8. *(i) A high-ability male prefers an informed audition if the evaluator's bias against female applicants is low; that is, $\beta \in \beta_L^B$. (ii) He also prefers an informed audition if the evaluator's bias against female applicants is high; that is, $\beta \in \beta_H^B$.*

Proof. (i) For $\beta \in \beta_L^B \equiv [0, \sqrt{p^2+16p}-p-2]$ and $\frac{1}{3} \leq p < 1$, the applicant's expected utility is strictly greater in an informed audition audition:

$$U(e_H, \eta_H, m) > U(e_H, \eta_H)$$
$$\Rightarrow 1 > \frac{2(1+p)+\beta}{4+2\beta}$$
$$\Rightarrow \beta > 2(p-1)$$

Given Proposition 2 and the assumption that evaluator's bias is non-negative, the inequality is always satisfied.

(ii) For $\beta \in \beta_H^B \equiv (\sqrt{17}-3, 2)$, the applicant's expected utility is also strictly greater in an informed audition, as shown in Figure 5.6. □

Intuitively, a high-ability male never wants to conceal his ability nor gender because his type always plays out in his favour. Under the condition that the evaluator's bias is strictly positive, even if the pool of applicants was predominantly of high ability, he would benefit from submitting his CV to identify himself as a male. Graphically, this means that, for any strictly positive bias, U_B will always be below U_I, even if the prior approaches one. This result contrasts with high-ability females whose preferences change to a blind audition as the pool of applicants becomes increasingly skilled.

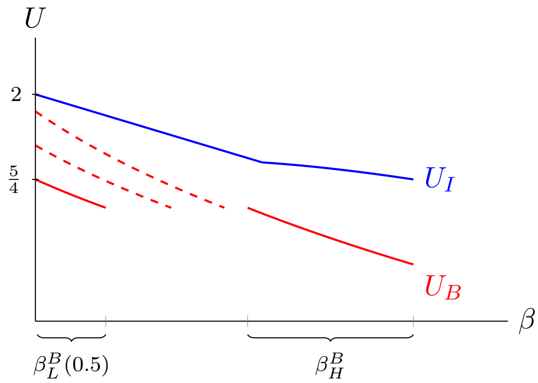

Figure 5.6: Expected Utility of High-ability Male in Blind and Informed Audition

5.2 Discussion of Results

Having discussed the preferences of all four types of applicants for different values of the prior and bias, I conclude that, if the evaluator's bias against female applicants is low, applicants' audition preferences differ along the ability dimension of their types; that is, high-ability applicants prefer an informed audition to identify themselves as talented. Low-ability applicants prefer a blind audition to pool with the talented.

In contrast, if the evaluator is highly biased, applicants' audition preferences differ along the gender dimension of their types; that is, males prefer an informed audition as the bias plays out in their favour. High-ability females prefer a blind audition as they rather trade off not being able to reveal their ability for not having to reveal their gender.

Not surprisingly, the evaluator's preferences align with the preferences of high-ability applicants if his bias against females is low. The evaluator wishes to screen applicants by their ability. In this case, he does so by requiring a CV which results in targeted effort incentives. High-ability applicants favour this policy as they want to distinguish themselves from low-ability applicants.

More surprisingly, the evaluator's preferences align with the preferences of high-ability females if his bias against them is large. In this case, to provide targeted effort incentives, the evaluator does not require a CV and holds the audition behind a curtain. This deters low-ability applicants and, as a side effect, plays out in favour of high-ability females.

Overall, the predictions of my model are qualitatively in line with the findings of Goldin and Rouse (2000) that "the switch to blind auditions can explain 30 percent of the increase in the proportion female among new hires and possibly 25 percent of the increase in the percentage female in the orchestras from 1970 to 1996" (p.738). Yet, my model sheds light on the underlying forces. First, if there is at most a low bias against female applicants, such that audition preferences differ along the ability dimension, the authors' observations are driven by the possibility of hiring low-ability females in a blind audition. In fact, the nonzero probability of hiring low-ability females more than compensates for the *reduced* probability of hiring high-ability females in a blind audition when the bias is low. To see this, note that the reduced probability is essentially the driver for why high-ability females prefer an informed audition as their effort level remains unchanged. In other words, for the same effort costs they are less likely to be hired in a blind audition. Second, if the bias against hiring female applicants is high, such that audition preferences differ along the gender dimension, their observations are entirely driven by the *increased* probability of hiring high-ability females in a blind audition, with no change for their low-ability counterparts. Analogous to the first case, the increased hiring probability essentially determines the preferences of high-ability females. In other words, for the same effort costs they are more likely to be hired in a blind audition.

5.3 Taste-based vs. Statistical Discrimination

Given my results, the model challenges existing explanations for the findings of Goldin and Rouse (2000) that are grounded in statistical discrimination. In particular, the main competing statistical discrimination model by Taylor and Yildirim (2011) endows the applicant only with a *one-dimensional* type θ that reflects her true productivity and is interpreted as either innate ability or experience. Furthermore, the applicant's type is drawn from some commonly known distribution function. Together, these assumptions imply that there is no room for the evaluator to bias irrelevant information, such as gender that may be contained in the applicant's CV, or for the evaluator's perception of ability to be misguided when screening applicants. However, drawing on the literature on cognitive biases, the widespread view of blind auditions as equal opportunity policy, the lawsuit of the African American bassist Art Davis, as well as the demand for unconventional entertainment formats and identity-withholding software (see Chapter 1), this is clearly a strong assumption.

A second aspect, in which the model by Taylor and Yildirim (2011) differs from the model laid out in this thesis, is the objective of the evaluator. In Taylor and Yildirim (2011), the evaluator wants to accept good and reject bad performances. Formally, he derives an exogenous benefit v or cost c from accepting a good or bad project, respectively. In my model, in contrast, the evaluator's objective is governed by the hiring rule in (3.1) or (4.1) in the third stage of an informed or blind audition, respectively. Either is derived from explicit structural assumptions on the production technology $f(q, n)$ according to which the applicant contributes to the orchestra if hired. While such explicit modelling poses the risk of a misspecified model, I offer a robustness check in Appendix A and I am able to identify the key mechanisms underlying the observations of Goldin and Rouse (2000) to make recommendations for policy. By specifying $f(q, n)$, I also make explicit that the applicant's marketable talent consist of both ability as a constant and performance quality, or equivalently effort, as a time-variant component. Taylor and Yildirim (2011) make the subtle assumption that the evaluator can never directly benefit from ability. This rules out revenue-generating reputation effects from hiring highly able star soloists[1] that have been shown to attract fans (Hart, 1973, p. 390) and additional single-ticket buyers (Kamakura and Schimmel, 2013).

Furthermore, my model is more realistic in the hiring policy: the evaluator always fills the position by either hiring the applicant or a random outside option at least temporarily, which is absent in Taylor and Yildirim (2011). Intuitively, in order to ensure a smooth running of the organisation, an evaluator is unlikely to leave a position unfilled. Concert schedules are planned far in advance, ticket sales tend to start months before the performances and orchestral parts need to be acquired from the publishers (Towse, 2010, chap. 8). Each part needs to be allocated to a musician before the rehearsals can begin. Given this inflexible nature of the performing arts, an unfilled position during the playing season would be very costly for the evaluator.

The strategically dynamic commitment benchmark in Taylor and Yildirim (2011) is comparable to the informed audition in this thesis. In particular, the authors assume that the evaluator is able to credibly commit to an acceptance threshold. Acting as a Stackelberg leader, the evaluator then maximises his expected payoff subject to the applicant's reaction function. In so doing, he faces a trade-off in providing effort incentives and efficiently using information: the evaluator is forced to reject performances that he knows to be of high quality with positive probability in order to provide effort incentives. Such an empirically unrealistic trade-off at the high end of the quality spectrum does not arise in my model. Yet, more similar to my model, applicants with ability *below* some threshold are screened: they drop out because they infer that the evaluator would never hire those types. The blind audition in Taylor and Yildirim (2011) is comparable to the one laid out in this thesis: the evaluator does not learn ability and, therefore, uses a uniform acceptance threshold. In my model, such a uniform threshold occurs when ability cannot be signalled; that is, in the pooling case, all applicants irrespective of gender and ability face the same hiring probabilities. Such empirically realistic signalling concerns à la Spence (1973) in the blind audition are absent in Taylor and Yildirim (2011).

The authors' main result is that the evaluator's audition preference depends on two dimensions of the environment: first, how able the applicant pool is and, second, how precise the signal of performance quality is. If the applicant pool is of mainly high ability,

[1]See also the economics literature on the phenomenon of superstars (Adler, 1985; Rosen, 1981).

the evaluator prefers a blind audition because assessing performance quality is relatively less important than providing effort incentives. If the applicant pool is of predominantly low ability, the evaluator prefers an informed audition because selecting high-quality performances is crucial. My model, in contrast, distinguishes whether, and how strongly, the evaluator is biased to make more nuanced predictions and derive recommendations for policy. If there is little bias, so that the evaluator is almost impartial, and the applicant pool is highly able, then a blind audition attracts low-ability applicants rather than providing targeted effort incentives. For an intermediate level of bias, my results are diametrically opposed to Taylor and Yildirim (2011): for a predominantly low-ability pool, a blind audition is preferred and vice versa. This is because the bias is already at a sufficiently high level to deter the less able from participating in the blind audition in the first place. For a high bias, the evaluator prefers a blind audition irrespective of the ability composition of the applicant pool.

Second, in Taylor and Yildirim (2011), if the signal of performance quality is very precise, he prefers a blind audition because the incremental information from observing ability is low. Conversely, if it is very imprecise, he prefers an informed audition because observing ability has a large incremental informational value. Due to Assumption 3, my model cannot shed light on the importance of uncertainty in the environment. Therefore, Chapter 6 relaxes this assumption through asymmetric uncertainty in the audition process. This enables me to provide a more nuanced comparison with the findings of Taylor and Yildirim (2011) in terms of signal precision.

6

Introducing Uncertainty

In this chapter, I consider the interaction between effort incentives and moral hazard by considering settings where effort does not uniquely determine performance quality. Assumption 3 shut down the moral hazard channel to focus on the interaction between the strength of the evaluator's bias and the information structure of the audition. In general, however, a key factor may be how strongly the applicant's performance quality on audition day correlates with her effort. Due to randomness in the audition environment, such as the room temperature or the humidity of the venue (Levinson, 2017), relying on a blindfolded performance may be a risky strategy for the evaluator.

To relax Assumption 3, I introduce *asymmetric* uncertainty into the audition process.

Assumption 4.

$$\Pr(q_H|e_H) = 1 \textit{ with } \Pr(q_H|e_L) = 0$$

and

$$\Pr(q_H|e_L) = \epsilon \textit{ with } \Pr(q_L|e_L) = 1 - \epsilon.$$

Performance quality is stochastically determined by the applicant's effort and one additional uncertainty parameter $\epsilon \in [0, \frac{1}{2})$. As ϵ gets close to zero, the extended model approaches the baseline model.

The asymmetry is motivated by the observation that high effort in the performing arts commonly takes the form of *deliberate practice* (Ericsson et al., 1993; Ericsson, 2006); that is, structured, intentional rehearsals with the goal to gradually build performance excellence. This may include learning as much as possible about performance-controlling factors (McPherson and Schubert, 2004). Musicians themselves often describe the process as one of overlearning and overpreparation: "[y]ou rehearse until you know what you are doing" (Brenda in Hays and Brown, 2006, p.97). In other words, deliberate practice removes the element of randomness, or luck, in the mapping from *high* effort to *high* performance quality because it allows "the performance to become sufficiently part of oneself such that the response becomes automatic, regardless of what happens" (Hays and Brown, 2006, p.99). Potential mechanisms for this one-to-one mapping are the prevention of performance anxiety, the ability to cope with pressure and the granting of self-confidence.

Conversely, low effort, interpreted as not engaging in deliberate practice, does not have a deterministic outcome. This is because perseverance is not the only route to performance excellence. In this case, randomness in the audition environment, or luck, can play a major role in the mapping: "heat, light, noise, as well as any other conditions that might interfere with obtaining a fair assessment of an individual's performance" (Castiglione, 1985, p.34). Alternatively, the applicant may have mastered the *inner game* and does not have to try so hard to achieve performance excellence nonetheless (Gallwey, 1974).

Assumption 4 may also capture the inherent subjectivity in performance evaluation; that is, the evaluator can identify who engaged in deliberate practice but his standards are too lenient to always identify at the margin those applicants who did not.

6.1 Informed Audition under Uncertainty

The hiring decision of the evaluator remains unchanged because performance quality is realised in the third stage. Given the unchanged ex-post hiring probabilities in (3.2), a risk-neutral applicant chooses effort to maximise her expected payoff. Substituting for the eight possible cases in the applicant's utility, taking into account the uncertainty in the mapping from low effort to performance quality, gives the following expected utilities:

$$U(e_L,\eta_L,f)=\frac{\epsilon}{2+\beta}w-\frac{1}{2} \qquad U(e_L,\eta_L,m)=\left[\frac{\beta(1-\epsilon)}{2+\beta}+\frac{(1+\beta)\epsilon}{2+\beta}\right]w-\frac{1}{2}$$

$$U(e_H,\eta_L,f)=\frac{1}{2+\beta}w-2 \qquad U(e_H,\eta_L,m)=\frac{1+\beta}{2+\beta}w-2$$

$$U(e_L,\eta_H,f)=\left[\frac{1-\epsilon}{2+\beta}+\frac{2\epsilon}{2+\beta}\right]w-\frac{1}{4} \qquad U(e_L,\eta_H,m)=\left[\frac{(1+\beta)(1-\epsilon)}{2+\beta}+\epsilon\right]w-\frac{1}{4}$$

$$U(e_H,\eta_H,f)=\frac{2}{2+\beta}w-1 \qquad U(e_H,\eta_H,m)=w-1$$

where $w\equiv\mathbb{E}[\bar{U}]=3-\frac{\beta}{2}$. Note that the effort cost is the same as in the baseline case. What changes is the ex-ante hiring probability that multiplies the applicant's wage when effort is low: it is a convex combination of low effort mapping into low performance quality, as well as the ϵ chance of attaining high performance quality as a consequence of randomness. Intuitively, such uncertainty, which could be interpreted as luck, makes exerting low effort relatively more attractive and alters the applicant's participation and incentive constraints. That, in turn, affects the applicant's optimal effort choice in the second stage.

Proposition 9. *For $\epsilon < \bar{\epsilon} \equiv \frac{1}{3}$, the effort responses of applicants under uncertainty partition the evaluator's bias $\beta \in [0, 2]$ into three regions: (i) For a low bias, $\beta \in \beta_L^I \equiv [0, \bar{\beta}_{LM})$, only high-ability applicants participate and exert high effort. (ii) For a moderate bias, $\beta \in \beta_M^I \equiv [\bar{\beta}_{LM}, \bar{\beta}_{HM}]$, low-ability males also participate and exert low effort. (iii) For a high bias, $\beta \in \beta_H^I \equiv (\bar{\beta}_{HM}, 2]$, all applicants except low-ability females participate and exert low effort.*

Proof. I consider the effect of uncertainty on incentive compatibility and participation constraints for all four types of applicants.

For low-ability females, a combination of high uncertainty and low bias makes participating and exerting low effort in an informed audition worthwhile. In particular, low-ability females exert low effort instead of dropping out if:

$$U(e_L, \eta_L, f) \geq 0 \Rightarrow 0 \leq \beta \leq \frac{2(3\epsilon - 1)}{\epsilon + 1}$$

Given that the bias is constrained to be non-negative, this results in the threshold $\bar{\beta}_{LF} \equiv \frac{2(3\epsilon-1)}{\epsilon+1}$ for $\epsilon \in [\frac{1}{3}, \frac{1}{2})$ and $\bar{\beta}_{LF} \equiv 0$ for $\epsilon \in [0, \frac{1}{3})$. At $\bar{\epsilon} \equiv \frac{1}{3}$, low effort is worthwhile if the evaluator is impartial. Beyond this critical value, the threshold bias for participation is increasing in ϵ:

$$\frac{d\bar{\beta}_{LF}}{d\epsilon} = \frac{8}{(\epsilon + 1)^2} > 0 \; \forall \; \epsilon \in [\frac{1}{3}, \frac{1}{2})$$

In other words, beyond a critical level of uncertainty, the evaluator can be increasingly biased against female applicants and still make low effort worthwhile (see Figure 6.1 where $\epsilon_0 = 0$).

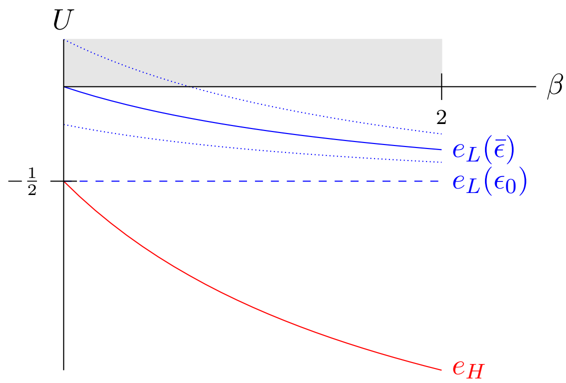

Figure 6.1: Effort Choice of Low-Ability Female in Informed Audition under Uncertainty

In contrast, uncertainty has an adverse effect on high-ability females. In particular, high-ability females exert low rather than high effort if:

$$U(e_L, \eta_H, f) \geq U(e_H, \eta_H, f) \Rightarrow \beta \geq \frac{6 - 12\epsilon}{5 - 2\epsilon}$$

Thus, $\bar{\beta}_{HF} \equiv \frac{6-12\epsilon}{5-2\epsilon}$ for $\epsilon \in [0, \frac{1}{2})$. For all levels of uncertainty, this threshold bias is decreasing in ϵ:

$$\frac{d\bar{\beta}_{HF}}{d\epsilon} = -\frac{48}{(5-2\epsilon)^2} < 0 \; \forall \; \epsilon \in [0, \frac{1}{2})$$

In fact, as $\epsilon \uparrow \frac{1}{2}$, the threshold bias approaches zero (see Figure 6.2). Intuitively, the more uncertain the environment, the lower the bias has to be to distort effort incentives for a highly able female.

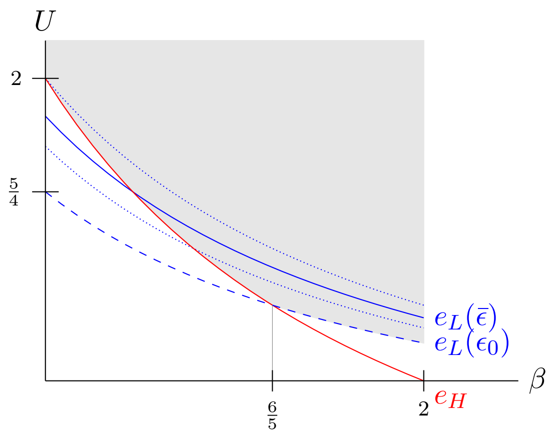

Figure 6.2: Effort Choice of High-Ability Female in Informed Audition under Uncertainty

For low-ability males, the threshold bias, beyond which participating and exerting low effort is worthwhile, falls with the introduction of uncertainty to the environment. In particular, low-ability males exert low effort instead of dropping out if:

$$U(e_L, \eta_L, m) \geq 0 \Rightarrow \beta \geq \frac{5-\epsilon-\sqrt{\epsilon^2+14\epsilon+17}}{2} \geq 0$$

Thus, $\bar{\beta}_{LM} \equiv \frac{5-\epsilon-\sqrt{\epsilon^2+14\epsilon+17}}{2}$ for $\epsilon \in [0, \frac{1}{3})$ and $\bar{\beta}_{LM} \equiv 0$ for $\epsilon \in [\frac{1}{3}, \frac{1}{2})$. For $\epsilon < \bar{\epsilon}$, this threshold bias is decreasing in ϵ:

$$\frac{d\bar{\beta}_{LM}}{d\epsilon} = -\frac{1}{2}\left[1 + \frac{\epsilon+7}{\sqrt{\epsilon^2+14\epsilon+17}}\right] < 0 \; \forall \; \epsilon \in [0, \frac{1}{3})$$

For $\epsilon \geq \bar{\epsilon}$, low-ability males are always induced to participate and exert low effort even if the evaluator is impartial (see Figure 6.3).

Uncertainty has an adverse effect on high-ability males. In particular, high-ability males exert low rather than high effort if:

$$U(e_L, \eta_H, m) \geq U(e_H, \eta_H, m) \Rightarrow \beta \geq \frac{6-12\epsilon}{5-2\epsilon}$$

Thus, $\bar{\beta}_{HM} \equiv \frac{6-12\epsilon}{5-2\epsilon}$ for $\epsilon \in [0, \frac{1}{2})$. Because this threshold bias coincides with the one derived for high-ability females, it follows immediately that $\bar{\beta}_{HM}$ is decreasing in ϵ. Similarly, as $\epsilon \uparrow \frac{1}{2}$, the threshold bias approaches zero (see Figure 6.4). □

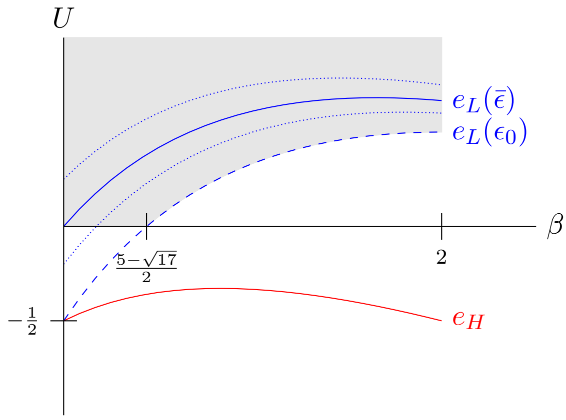

Figure 6.3: Effort Choice of Low-Ability Male in Informed Audition under Uncertainty

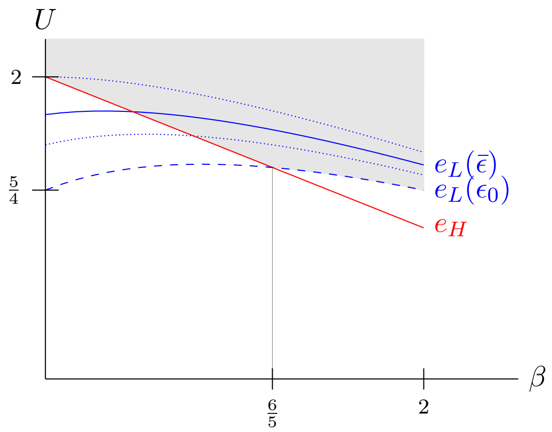

Figure 6.4: Effort Choice of High-Ability Male in Informed Audition under Uncertainty

To get some intuition behind Proposition 9; that is, how the applicants' effort responses under rising uncertainty partition the evaluator's bias $\beta \in [0, 2]$, fix $\epsilon_0 = 0$, $\epsilon_1 = \frac{1}{20}$, $\epsilon_2 = \frac{1}{10}$, $\epsilon_3 = \frac{1}{5}$ and $\bar{\epsilon} = \frac{1}{3}$. First, Figure 6.5 illustrates that the interval β_H^I is expanding to the left as the degree of uncertainty increases. Second, an increase in uncertainty causes the interval β_L^I to shrink at a faster rate than the interval β_M^I. As a result, $\bar{\epsilon} \equiv \frac{1}{3}$ marks the threshold beyond which β_L^I vanishes.

These two observations are intuitive: the more uncertain the environment, the lower the bias of the evaluator has to be to distort effort incentives for the highly able. This is exacerbated by the fact that uncertainty has an attracting effect on low-ability applicants, where $\bar{\epsilon} \equiv \frac{1}{3}$ is a threshold that determines which gender benefits from rising uncertainty: as the degree of uncertainty approaches the threshold from below, low-ability males benefit and participate for a greater range of biases, with no change for low-ability females. Beyond this threshold, as the degree of uncertainty approaches its supremum, low-ability females benefit and, subsequently, participate more often.

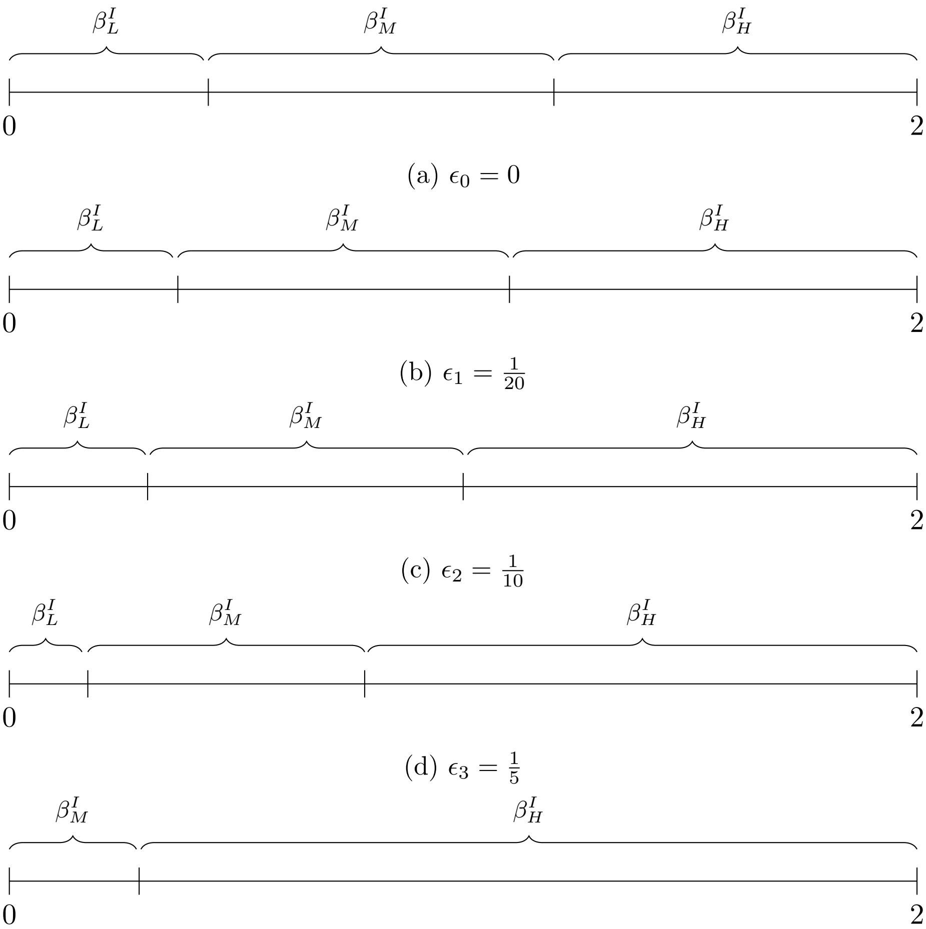

(a) $\epsilon_0 = 0$

(b) $\epsilon_1 = \frac{1}{20}$

(c) $\epsilon_2 = \frac{1}{10}$

(d) $\epsilon_3 = \frac{1}{5}$

(e) $\bar{\epsilon} = \frac{1}{3}$

Figure 6.5: Partition of Evaluator's Bias for Different Degrees of Uncertainty

6.1.1 Expected Utility of Evaluator in Stage 1 under Uncertainty

In what follows, I assume $\epsilon < \bar{\epsilon}$. While this restriction is not without loss of generality, it ensures that low-ability females never participate and that the results are comparable to those derived under certainty. Given $\epsilon < \bar{\epsilon}$, Proposition 9 holds and the applicants' effort responses partition the evaluator's bias $\beta \in [0, 2]$ into three intervals: β_L^I, β_M^I and β_H^I.

Suppose the evaluator's bias against female applicants is low; that is, $\beta \in \beta_L^I$ where β_L^I is defined as in Proposition 9(i). Then, the evaluator's expected net utility Π_I is not affected by uncertainty and is given by (3.3). This is because the bias-uncertainty combination is sufficiently low to provide targeted effort incentives.

Suppose the evaluator's bias against female applicants is moderate; that is, $\beta \in \beta_M^I$ where β_M^I is defined as in Proposition 9(ii). Under uncertainty, the evaluator's expected net utility Π_I modifies to:

$$\mathbb{E}[\Pi_I|\beta^I_M] = \frac{1-p}{2}\left[(1-\epsilon)\frac{\beta^I_M}{2+\beta^I_M}(\frac{\beta^I_M}{2}-1)+\epsilon\frac{1+\beta^I_M}{2+\beta^I_M}(\frac{\beta^I_M}{2})\right]$$
$$+\frac{p}{2}\left[\frac{2}{2+\beta^I_M}(1-\frac{\beta^I_M}{2})+(1+\frac{\beta^I_M}{2})\right] \gtrless 0$$

Suppose the evaluator's bias against female applicants is high; that is, $\beta \in \beta^I_H$ where β^I_H is defined as in Proposition 9(iii). Under uncertainty, the evaluator's expected net utility Π_I modifies to:

$$\mathbb{E}[\Pi_I|\beta^I_H] = \frac{1-p}{2}\left[(1-\epsilon)\frac{\beta^I_H}{2+\beta^I_H}(\frac{\beta^I_H}{2}-1)+\epsilon\frac{1+\beta^I_H}{2+\beta^I_H}(\frac{\beta^I_H}{2})\right]$$
$$+\frac{p}{2}\left[(1-\epsilon)\frac{1}{2+\beta^I_H}(-\frac{\beta^I_H}{2})+\epsilon\frac{2}{2+\beta^I_H}(1-\frac{\beta^I_H}{2})\right.$$
$$\left.+(1-\epsilon)\frac{1+\beta^I_H}{2+\beta^I_H}(\frac{\beta^I_H}{2})+\epsilon(1+\frac{\beta^I_H}{2})\right] \gtrless 0$$

6.2 Blind Audition under Uncertainty

I consider the effect of uncertainty on the pooling and separating equilibrium analysed in Chapter 4. A key question for policy is whether the separating equilibrium: (i) continues to exist and (ii) continues to be more profitable for the evaluator. This would ensure that my recommendations generalise to an uncertain environment and that a strategy of commitment to no information can be beneficial if the evaluator knows that he would otherwise not be impartial.

6.2.1 Pooling Equilibrium under Uncertainty

Suppose the evaluator's bias and the degree of uncertainty are sufficiently low to induce all applicants to exert high effort. Given the asymmetric uncertainty in Assumption 4, low performance quality continues to be off the equilibrium path. As in Chapter 4, I assume that, when observing a low-quality performance off the equilibrium path, the evaluator believes the applicant to be of low ability. Upon observing a low-quality performance, the evaluator can, therefore, infer the applicant's ability. Upon observing a high-quality performance, the evaluator cannot infer the applicant's ability and holds a belief that is equal to his prior; that is, $\Pr(\eta_H) = p \in (0,1)$. As a result, the hiring decision of the evaluator in the third stage remains unchanged with ex-post hiring probabilities given by (4.2). Substituting for the two possible cases in the applicant's utility, taking into account the uncertainty in the mapping from low effort to performance quality, gives the following expected utilities:

$$U(e_L,\eta_L) = \left[\frac{\beta}{4+2\beta}(1-\epsilon) + \frac{(2(1+p)+\beta)}{4+2\beta}\epsilon\right]w - \frac{1}{2}$$
$$U(e_H,\eta_L) = \frac{2(1+p)+\beta}{4+2\beta}w - 2$$
$$U(e_L,\eta_H) = \left[\frac{\beta}{4+2\beta}(1-\epsilon) + \frac{(2(1+p)+\beta)}{4+2\beta}\epsilon\right]w - \frac{1}{4}$$
$$U(e_H,\eta_H) = \frac{2(1+p)+\beta}{4+2\beta}w - 1$$

where $w \equiv \mathbb{E}[\bar{U}] = 3 - \frac{\beta}{2}$. I now show that under uncertainty p and ϵ usually determine jointly whether a high-effort pooling equilibrium can be supported and that p and ϵ place an upper bound on the evaluator's bias in any such pooling equilibrium.

Proposition 10. *(i) For $p < \frac{1}{3}$, no high-effort pooling equilibrium exists. (ii) For $p = \frac{1}{3}$, a high-effort pooling equilibrium exists if the evaluator is unbiased and $\epsilon \leq \frac{1}{4}$. (iii) For $\frac{1}{3} < p < 1$, a high-effort pooling equilibrium exists if $\beta \in \beta_L^B \equiv [0, \bar{\beta}_{pool}]$ and $\bar{\beta}_{pool} \geq 0$.*

Proof. (i) First, note that $U(e_H,\eta_L)$ is independent of ϵ. Therefore, by Proposition 2(i), for $p < \frac{1}{3}$, pooling on high effort is never possible.

(ii) For $p = \frac{1}{3}$, $U(e_H,\eta_L) = 0$ at $\beta = 0$. By Proposition 2(i), $U(e_H,\eta_L) < 0$ for all $\beta \in (0,2]$. This implies that the low-ability applicant is indifferent between exerting high effort and not participating only if the evaluator is unbiased. Under uncertainty, I also need to ensure incentive compatibility. Because it is optimal for high-ability applicants to exert high effort whenever it is optimal for low-ability applicants to exert high effort, this constraint is given by:

$$U(e_H,\eta_L)\Big|_{\beta=0,p=\frac{1}{3}} \geq U(e_L,\eta_L)\Big|_{\beta=0,p=\frac{1}{3}} \Rightarrow \epsilon \leq \frac{1}{4}$$

Thus, pooling on high effort is possible at $p = \frac{1}{3}$ if $\beta = 0$ and $\epsilon \leq \frac{1}{4}$.

(iii) For $\frac{1}{3} < p < 1$, for low degrees of uncertainty, the upper bound on β for which a pooling equilibrium can be supported is determined by the indifference condition:

$$U(e_H,\eta_L) = \frac{2(1+p)+\beta}{4+2}(3-\frac{\beta}{2}) - 2 = 0$$

Given that $\beta \in [0,2]$, the above equation is uniquely solved by $\beta(p) \equiv \sqrt{p^2+16p} - p - 2$.

For larger degrees of uncertainty, the upper bound on β for which a pooling equilibrium can be supported is determined by the indifference condition:

$$U(e_H,\eta_L) = U(e_L,\eta_L)$$

Given that $\beta \in [0, 2]$, $\epsilon \in [0, \frac{1}{2})$ and $p \in (0, 1)$, the above equation is uniquely solved by $\beta(p, \epsilon) \equiv \frac{6(p\epsilon - p + \epsilon)}{p\epsilon - p + \epsilon - 4}$. Therefore, the upper bound on a high-effort pooling equilibrium is given by $\bar{\beta}_{pool} \equiv \min\{\beta(p), \beta(p, \epsilon)\}$, provided that $\bar{\beta}_{pool}$ is non-negative. Otherwise, no high-effort pooling equilibrium exists. □

Proposition 10(iii) implies that pooling on high effort is possible if the prior that the applicant is of high ability is sufficiently large and the evaluator's bias as well as the degree of uncertainty is not too extreme. In particular, for a fixed prior $p > \frac{1}{3}$, once a critical level of uncertainty is reached, the range of biases for which pooling on high effort is an equilibrium is decreasing in uncertainty. This is driven by the fact that low effort becomes relatively more attractive. To see this intuitively, fix the prior at $p = \frac{1}{2}$. An increase in uncertainty then graphically implies an upward shift of the low-effort curve for both low- and high-ability applicants in (β, U)-space (see Figure 6.6 and Figure 6.7). In this particular example, the critical level of uncertainty is approximately 0.16.

Conversely, for a fixed degree of uncertainty $\epsilon > 0$, once a critical level of bias is reached, the range of biases for which pooling on high effort is an equilibrium outcome is increasing in the prior. Intuitively, while the ability composition of the applicant pool affects the applicant's utility under either effort choice, the effect of an increasingly able applicant pool is less pronounced when she chooses to exert low effort. This is because the effect of a change in the ability composition of the applicant pool is scaled down by the uncertainty that arises from exerting low effort. Therefore, this comparative statics result is robust to the introduction of uncertainty; that is, an increasingly able applicant pool makes it easier for for the fewer low-ability to hide behind the more-and-more high-ability applicants.

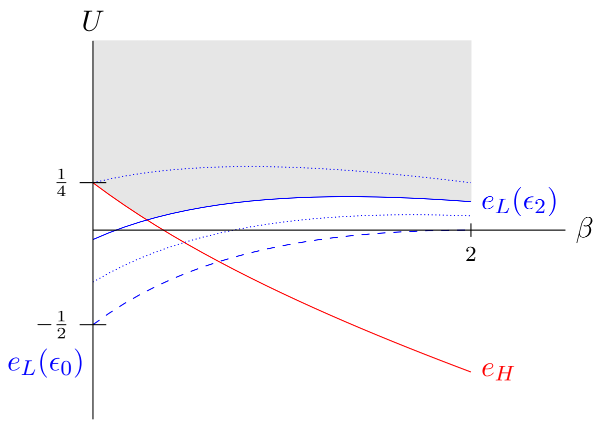

Figure 6.6: Effort Choice of Low-Ability Applicant in Pooling Blind Audition under Uncertainty for $\Pr(\eta_H) = \frac{1}{2}$

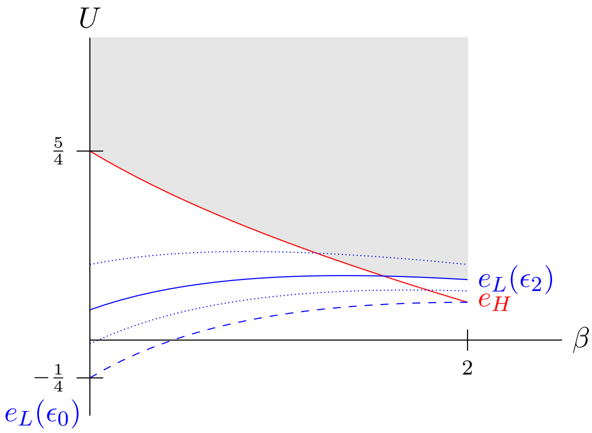

Figure 6.7: Effort Choice of High-Ability Applicant in Pooling Blind Audition under Uncertainty for $\Pr(\eta_H) = \frac{1}{2}$

6.2.1.1 Expected Utility of Evaluator in Stage 1 under Uncertainty

If a high-effort pooling equilibrium under uncertainty exists, the evaluator's expected net utility Π_B is not affected by uncertainty and is given by (4.3).

6.2.2 Fully Separating Equilibrium under (Un)certainty

Suppose the evaluator's bias is sufficiently high and the degree of uncertainty is sufficiently low to induce high-ability applicants to exert high effort and low-ability applicants to drop out. Under uncertainty, this qualification, which I refer to as *full separation*, is important: it ensures that the evaluator has degenerate posterior beliefs at the information set q_H. In other words, upon observing a high-quality performance, the evaluator is certain to face a highly able applicant, and male or female with equal probability. If, in contrast, low-ability applicants were to exert low effort, the evaluator could not be certain about the applicant's ability at q_H due to the ϵ probability that low effort results in a high-quality performance. In this case, I would need to use Bayes' rule in deriving $\mathbb{E}[V|q_H]$ and this would alter the hiring probability at q_H (see Section 6.2.3).

As in the pooling equilibrium, the evaluator believes the applicant to be of low ability when observing a low performance quality off the equilibrium path. Under the above qualification, the hiring decision of the evaluator in the third stage remains unchanged with ex-post hiring probabilities given by (4.4). Substituting for the two possible cases in the applicant's utility, taking into account the uncertainty in the mapping from low effort to performance quality, gives the following expected utilities:

$$
\begin{aligned}
U(e_L,\eta_L) &= \left[\frac{\beta}{4+2\beta}(1-\epsilon)+\frac{4+\beta}{4+2\beta}\epsilon\right]w-\frac{1}{2}\\
U(e_H,\eta_L) &= \frac{4+\beta}{4+2\beta}w-2\\
U(e_L,\eta_H) &= \left[\frac{\beta}{4+2\beta}(1-\epsilon)+\frac{4+\beta}{4+2\beta}\epsilon\right]w-\frac{1}{4}\\
U(e_H,\eta_H) &= \frac{4+\beta}{4+2\beta}w-1
\end{aligned}
$$

where $w \equiv \mathbb{E}[\bar{U}] = 3 - \frac{\beta}{2}$. Because the prior that the applicant is of high ability does not enter the applicant's expected utility, I can compare participation and incentive constraints as in an informed audition to characterise equilibrium in terms of the bias parameter and the degree of uncertainty.

Figure 6.8 reveals that even the slightest degree of uncertainty in the environment makes participating and exerting low effort for low-ability applicants worthwhile. More precisely, beyond a *negligible* degree of uncertainty $\hat{\epsilon} \equiv \frac{13}{16} - \frac{3\sqrt{17}}{16} \approx \frac{1}{25}$ (see solid blue line for $\hat{\epsilon}$), there is no range of biases that can support a separating equilibrium in which low-ability applicants do not participate. This frailty of the fully separating equilibrium to the introduction of uncertainty stems from the fact that the payoff from exerting low effort for the less able is just below zero. As a result, the *deterrence effect* of the evaluator's beliefs off the equilibrium path is voided by the slightest chance of attaining high performance quality as a consequence of randomness or by having a lucky day.

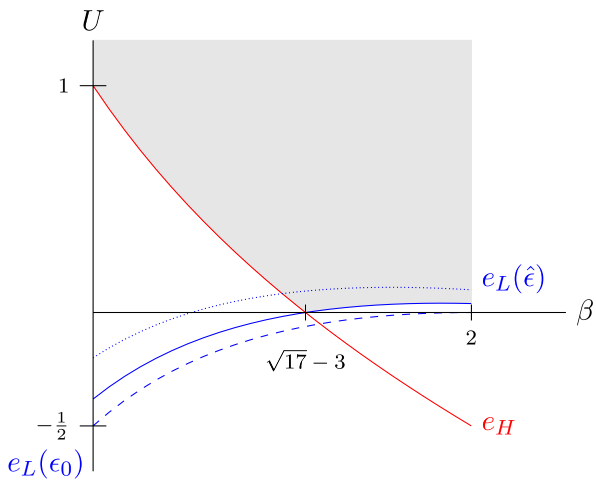

Figure 6.8: Effort Choice of Low-Ability Applicant in Fully Separating Blind Audition under (Un)certainty

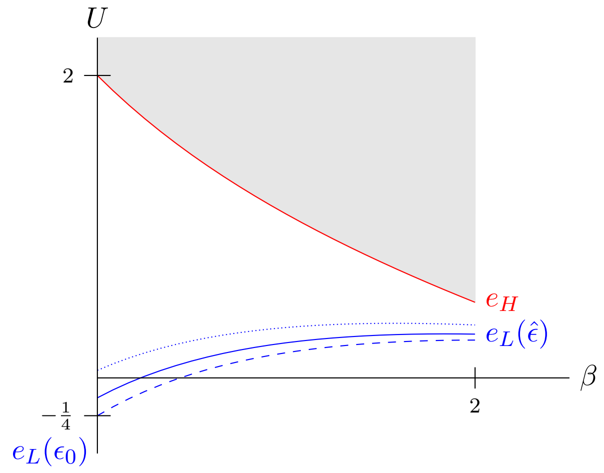

Figure 6.9: Effort Choice of High-Ability Applicant in Fully Separating Blind Audition under (Un)certainty

6.2.2.1 Expected Utility of Evaluator in Stage 1 under (Un)certainty

For the evaluator's belief that only high-ability applicants participate and exert high effort to be consistent, he needs to have a sufficiently high bias β_H^B and the degree of uncertainty has to be negligible. In this case, low-ability applicants do not participate and q_L is not reached. The evaluator's expected net utility Π_B is not affected by the negligible degree of uncertainty and is given by (4.5).

6.2.3 Partially Separating Equilibrium under Uncertainty

The frailty of the fully separating equilibrium to the introduction of uncertainty suggests that there exists a less extreme separating equilibrium for a range of bias-prior-uncertainty combinations, in which highly able applicants exert high effort and the less able exert low effort. I refer to this outcome as *partial separation*. Note that partial separation was not an equilibrium under certainty: I specified the evaluator's beliefs off the equilibrium path to be such that he believes the applicant to be of low ability with probability one and I did not restrict ex-ante whether the less able, in fact, participate or not. However, for $\beta < 2$, the outside option turned out to dominate low effort for the less able and I reverted to two refinements to show robustness.

Suppose the evaluator's bias and uncertainty are sufficiently high to induce high-ability applicants to exert high effort and low-ability applicants to exert low effort. When observing a low-quality performance, the evaluator is certain that the applicant is of low ability. By Bayes' rule, when observing a high-quality performance, the evaluator believes the applicant to be of high ability with probability:

$$\Pr(\eta_H|q_H) = \frac{p}{p+\epsilon(1-p)} \equiv \gamma$$

and of low ability with complementary probability. At q_H, the evaluator, therefore, expects gross utility:

$$\begin{aligned}
\mathbb{E}[V|q_H] &= q_H + \mathbb{E}[\eta|q_H] - \frac{\beta}{2} = 2 + \left[\frac{\epsilon(1-p)}{p+\epsilon(1-p)} + \frac{2p}{p+\epsilon(1-p)}\right] - \frac{\beta}{2} \\
&= 2 + \left[1 + \frac{p}{p+\epsilon(1-p)}\right] - \frac{\beta}{2} \\
&= 3 + \gamma - \frac{\beta}{2}
\end{aligned}$$

from hiring the applicant. Therefore, the hiring probability at q_H modifies to:

$$\Pr(h|q_H) = \Pr(\bar{U} \leq 3 + \gamma - \frac{\beta}{2}) = \frac{2(1+\gamma)+\beta}{4+2\beta}$$

which is strictly smaller than in the fully separating equilibrium for all $\epsilon > 0$. This is intuitive: under partial separation, the evaluator is more cautious in his hiring policy when observing a high-quality performance due to the chance that the applicant behind the curtain is only of low ability. The hiring probability at q_L remains unchanged.

Substituting for the two possible cases in the applicant's utility, taking into account the uncertainty in the mapping from low effort to performance quality, gives the following expected utilities:

$$\begin{aligned}
U(e_L, \eta_L) &= \left[\frac{\beta}{4+2\beta}(1-\epsilon) + \frac{2(1+\gamma)+\beta}{4+2\beta}\epsilon\right] w - \frac{1}{2} \\
U(e_H, \eta_L) &= \frac{2(1+\gamma)+\beta}{4+2\beta} w - 2 \\
U(e_L, \eta_H) &= \left[\frac{\beta}{4+2\beta}(1-\epsilon) + \frac{2(1+\gamma)+\beta}{4+2\beta}\epsilon\right] w - \frac{1}{4} \\
U(e_H, \eta_H) &= \frac{2(1+\gamma)+\beta}{4+2\beta} w - 1
\end{aligned}$$

where $w \equiv \mathbb{E}[\bar{U}] = 3 - \frac{\beta}{2}$.

Proposition 11. *For $\beta \in \beta_H^B \equiv (\underline{\beta}_{sep}, \bar{\beta}_{sep}]$, a partially separating equilibrium exists in which high-ability applicants exert high effort and low-ability applicants exert low effort.*

Proof. First, note that low-ability applicants must find it optimal to exert low effort. This gives the lower bound $\underline{\beta}_{sep}$ for which a partially separating equilibrium can be supported. For low degrees of uncertainty, the lower bound is determined by the indifference condition:

$$U(e_L, \eta_L) = 0$$

Given that $\beta \in [0, 2]$, $\epsilon \in [0, \frac{1}{2})$ and $p \in (0, 1)$, the above equation is uniquely solved by $\underline{\beta}_{sep'} \equiv \frac{\sqrt{p^2\epsilon^4+4p^2\epsilon^3-20p^2\epsilon^2+16p^2\epsilon-2p\epsilon^4-12p\epsilon^3+24p\epsilon^2+\epsilon^4+8\epsilon^3}-p\epsilon^2+4p\epsilon-2p+\epsilon^2-2\epsilon}{p\epsilon-p-\epsilon}$. For larger degrees of uncertainty, the lower bound is determined by the indifference condition:

$$U(e_L, \eta_L) = U(e_H, \eta_L)$$

Given that $\beta \in [0, 2]$, $\epsilon \in [0, \frac{1}{2})$ and $p \in (0, 1)$, the above equation is uniquely solved by $\underline{\beta}_{sep''} \equiv \frac{6(p\epsilon^2-2p\epsilon+p-\epsilon^2)}{p\epsilon^2-6p\epsilon+5p-\epsilon^2+4\epsilon}$. Therefore, the lower bound of the partially separating equilibrium is given by $\underline{\beta}_{sep} \equiv \max\{0, \underline{\beta}_{sep'}, \underline{\beta}_{sep''}\}$.

Second, note that high-ability applicants must find it optimal to exert high effort. This gives the upper bound $\bar{\beta}_{sep}$ for which a partially separating equilibrium can be supported. For high degrees of uncertainty, the upper bound is determined by the indifference condition:

$$U(e_L, \eta_H) = U(e_H, \eta_H)$$

Given that $\beta \in [0, 2]$, $\epsilon \in [0, \frac{1}{2})$ and $p \in (0, 1)$, the above equation is uniquely solved by $\bar{\beta}_{sep'} \equiv \frac{6(2p\epsilon^2-5p\epsilon+3p-2\epsilon^2+\epsilon)}{2p\epsilon^2-9p\epsilon+7p-2\epsilon^2+5\epsilon}$. Therefore, the upper bound of the partially separating equilibrium is given by $\bar{\beta}_{sep} \equiv \min\{2, \bar{\beta}_{sep'}\}$.

$\square$

Proposition 11 implies that, for a fixed prior, once a negligible degree of uncertainty is reached, partial separation is an equilibrium for a significant range of biases[1]. As the prior enters the applicant's expected utility via the conditional probability that she is of high ability given a high-quality performance, it is instructive to fix the prior at $p = \frac{1}{2}$ to get some intuition. Note also that the degree of uncertainty affects the applicant's expected utility from low and high effort. As a result, both $U(e_H, \cdot)$ and $U(e_L, \cdot)$ shift in response to changes in ϵ in (β, U)-space, making comparative statics difficult. To make some progress, Figure 6.10 and Figure 6.11 show applicants' effort responses for two levels of uncertainty relative to the dashed baseline $\epsilon_0 = 0$. For $\epsilon_2 = \frac{1}{10}$ (solid lines), the range of biases for which partial separation is an equilibrium is given by $\beta_H^B = [\frac{158}{173}, 2]$. For $\epsilon_4 = \frac{3}{10}$ (dotted lines), the range of biases for which partial separation is an equilibrium is given by $\beta_H^B = [\frac{186}{551}, \frac{144}{89}]$.

An increase in uncertainty from ϵ_2 to ϵ_4 makes exerting low effort relatively more attractive for both high- and low-ability applicants. As a result, it reduces the lower bound of the partially separating equilibrium via the effort responses of the less able. Beyond a critical level, an increase in uncertainty also reduces the upper bound of the partially separating equilibrium via the effort responses of the highly able. For $p = \frac{1}{2}$, this critical level of uncertainty is approximately 0.19.

[1]At $\epsilon_0 = 0$, low-ability applicants are indifferent between exerting low effort and not participating if the evaluator's bias is maximal, $\beta = 2$. Thus, I defined $\beta_H^B \equiv (\sqrt{17} - 3, 2)$ in Chapter 4 to support the fully separating equilibrium and neglected the possibility of partial separation at $\beta = 2$.

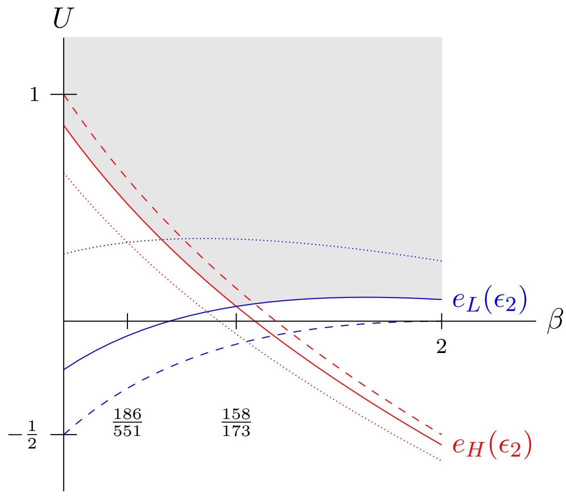

Figure 6.10: Effort Choice of Low-Ability Applicant in Partially Separating Blind Audition under Uncertainty

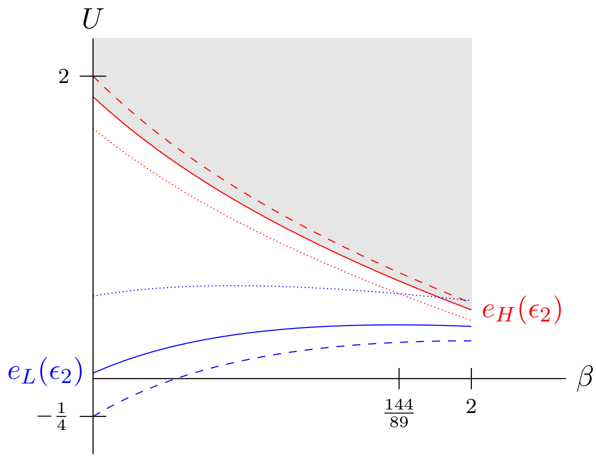

Figure 6.11: Effort Choice of High-Ability Applicant in Partially Separating Blind Audition under Uncertainty

6.2.3.1 Expected Utility of Evaluator in Stage 1 under Uncertainty

For the evaluator's belief that high-ability applicants exert high effort and low-ability applicants exert low effort to be consistent, he needs to have a sufficiently high bias β_H^B and

$\epsilon > 0$. The evaluator's expected net utility Π_B is:

$$\mathbb{E}[\Pi_B|\beta_H^B] =$$
$$(1-p)\left[\frac{\beta_H^B}{4+2\beta_H^B}(1-\epsilon)\left(2-\frac{\beta_H^B}{2}-(3-\frac{\beta_H^B}{2})\right)+\frac{2(1+\gamma)+\beta_H^B}{4+2\beta_H^B}\epsilon\left(3-\frac{\beta_H^B}{2}-(3-\frac{\beta_H^B}{2})\right)\right]$$
$$+p\left[\frac{2(1+\gamma)+\beta_H^B}{4+2\beta_H^B}\left(4-\frac{\beta_H^B}{2}-(3-\frac{\beta_H^B}{2})\right)\right]$$
$$=(1-p)\left[-\frac{\beta_H^B}{4+2\beta_H^B}(1-\epsilon)\right]+p\left[\frac{2(1+\gamma)+\beta_H^B}{4+2\beta_H^B}\right]$$

6.3 Comparison of Blind and Informed Audition under Uncertainty

Having discussed the effect of uncertainty on the evaluator's net utility in both forms of audition, I can conclude under which conditions on the bias parameter β, the prior p, and the degree of uncertainty ϵ my results on the evaluator's preferences continue to hold.

First, note that, as full separation under uncertainty is not possible, blind auditions do not have the same fully fledged effect as under certainty; that is, blind auditions cannot provide targeted effort incentives when the evaluator is highly biased. This dampens the evaluator's expected net utility. For a fixed prior, compare, for example, the evaluator's full separation net utility when $\epsilon_0 = 0$ in Figure 6.12a with his partial separation net utility when $\epsilon_2 = \frac{1}{10}$ in Figure 6.12b: while a highly biased evaluator still finds blind auditions more profitable than informed auditions for this ability composition of the applicant pool and this degree of uncertainty, the expected net gain is smaller.

Second, under uncertainty, blind auditions are not guaranteed to be more profitable than informed auditions when the evaluator is highly biased (see Figure 6.13). However, for any degree of uncertainty, for the range of biases for which partial separation can be supported, blind auditions are still guaranteed to be more profitable for a sufficiently high prior approximately greater than 0.46. Intuitively, while the less able participate in the blind audition, the highly biased evaluator can, at least, provide effort incentives for roughly the majority of the highly able.

Third, full separation insured the highly biased evaluator against negative net utility in a blind audition. Under uncertainty and partial separation, the evaluator's net utility may be negative for very low priors. In fact, there are bias-prior-uncertainty combinations for which neither blind nor informed auditions are profitable. For $\epsilon_2 = \frac{1}{10}$, these combinations are the grey-shaded areas in Figure 6.13.

To gain some intuition behind the evaluator's distribution of preferences, suppose that, if the degree of uncertainty ϵ_2 is moderate, the evaluator's bias and the prior on high ability are drawn uniformly from their respective supports. In other words, all points in Figure 6.13 are equally likely before the three-stage game under uncertainty is played. Therefore, I can use integration to calculate the ex-ante percentage with which a blind audition (red area) is preferred relative to an informed audition (blue area), or neither (grey area). In so doing, I discount the white areas for which no comparison is possible. The evaluator prefers a blind audition approximately fifty percent and an informed audition approximately forty-five percent of the time. Neither audition is preferred approximately five percent of the time. First, this highlights the sizeable risk of market failure under uncertainty; that is, the evaluator might not want to hold an audition at all. Second, this highlights that an informed audition gains only about nine percentage points while a blind audition loses about fourteen percentage points in the move from a certain to a moderately uncertain environment. In other words, rather than a redistribution of preferences, the evaluator is worse off because the uncertainty reduces the attractiveness of a blind audition by more than it increases the attractiveness of an informed audition. At a macroeconomic level, the moderate uncertainty combined with this simple form of heterogeneity in the bias across evaluators as well as in the ability composition of applicant pools predicts the co-existence of blind and informed auditions in more equal proportions.

6.3.1 Taste-based vs. Statistical Discrimination under Uncertainty

Given my results, the model challenges the competing statistical discrimination model also under uncertainty. In Taylor and Yildirim (2011), if the signal of performance quality is very precise, the evaluator prefers a blind audition because the incremental information from observing ability is low. Conversely, if it is very imprecise, he prefers an informed audition because observing ability has a large incremental informational value.

My model, in contrast, distinguishes whether, and how strongly, the evaluator is biased to make more nuanced predictions about the effect of uncertainty on the evaluator's preferences and derive recommendations for policy. First, my results contrast with Taylor and Yildirim (2011) if the evaluator is almost impartial: informed auditions are preferred irrespective of the degree of uncertainty in the environment.

Second, my results align with Taylor and Yildirim (2011) if the evaluator is highly biased: if there is a *negligible* degree of uncertainty, such that full separation is possible, blind auditions attract only the highly able, allowing the evaluator to infer ability. In this case, the informational gain from requiring a CV is small; if anything, it would distort effort incentives. If there is a *non-negligible* degree of uncertainty, the ability composition

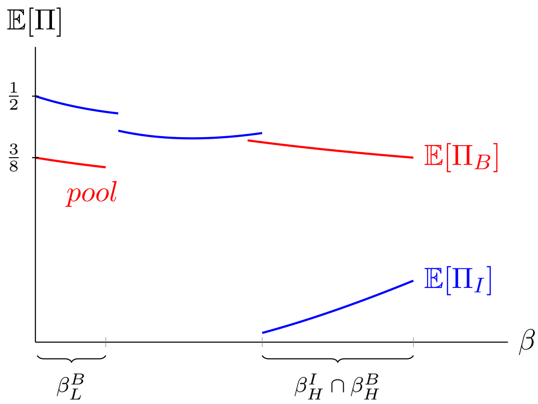

(a) $\Pr(\eta_H) = \frac{1}{2}$ and $\epsilon_0 = 0$

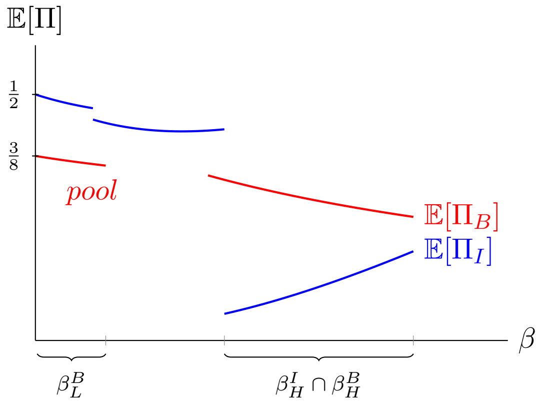

(b) $\Pr(\eta_H) = \frac{1}{2}$ and $\epsilon_2 = \frac{1}{10}$

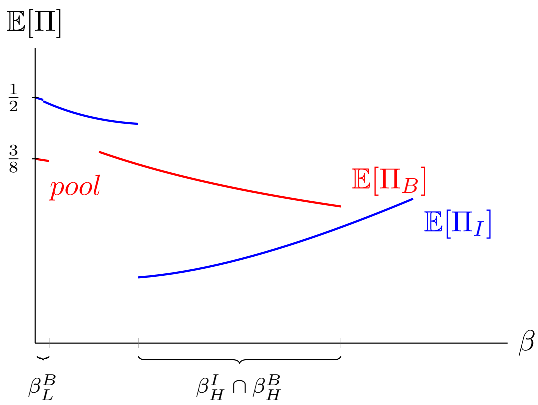

(c) $\Pr(\eta_H) = \frac{1}{2}$ and $\epsilon_4 = \frac{3}{10}$

Figure 6.12: Evaluator's Expected Net Utility in Blind and Informed Audition under Uncertainty

of the applicant pool matters for the evaluator's preferences. For a *very low* prior[2], it is paramount for the evaluator to include ability in his hiring policy. While blind auditions

[2] A *very low* prior is implicitly defined by the degree of uncertainty and a high bias such that an informed audition is preferable. In Figure 6.13, for example, for $\epsilon_2 = \frac{1}{10}$ and $\beta \in \beta_H^I$, all priors in the blue region satisfy this definition.

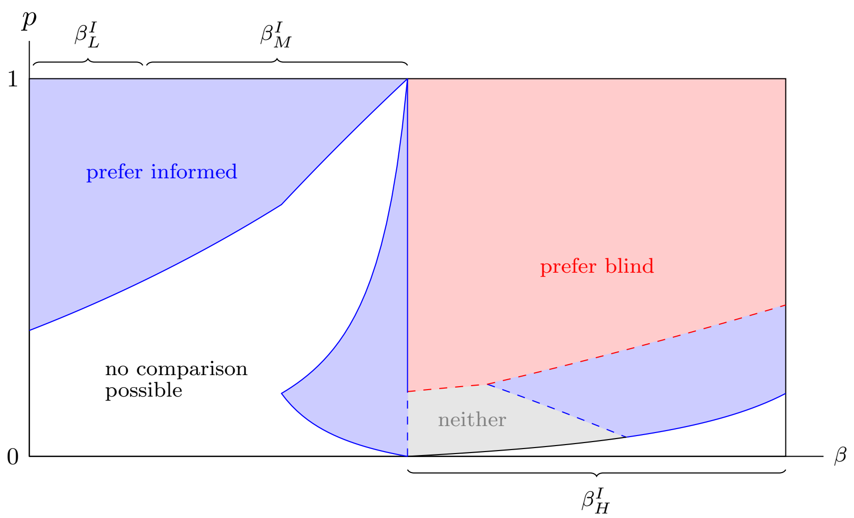

Figure 6.13: Evaluator's Preferences over Auditions for Different Biases and Priors for $\epsilon_2 = \frac{1}{10}$

provide better effort incentives for the highly able, these types are rather infrequent in the audition, making informed auditions preferable. For a *sufficiently high* prior, high effort incentives for the highly able are paramount, making blind auditions preferable.

6.3.2 Policy Recommendations under Uncertainty

If uncertainty is an inherent feature of the environment and cannot be alleviated, then policy should ensure that the applicant pool is of sufficiently high ability. Examples could be early childhood interventions to identify talent and to build performance excellence gradually via deliberate practice and, thereafter, merit scholarships for formal training to secure a steady influx of highly able musicians into the applicant pool. Direct interventions could target artistic poverty and the necessity for multi-jobbing, often in teaching, to earn a subsistence income as main barriers to human capital investment later in life. One could ask whether "auditioning, networking, keeping fit, practising and unpaid rehearsing [should] be counted as part of the artists' labo[u]r supply" (Towse, 1996, p.102) to reward human capital investment.

Under uncertainty, such ability-targeting policy is beneficial for two reasons: first, it guarantees the profitability of the partially separating blind audition when the evaluator is highly biased. Second, it guarantees that blind auditions still do better than informed auditions. This may be desirable given equity concerns and equality of opportunity along the gender dimension of applicants.

Intuitively, under uncertainty, an applicant pool of sufficiently high ability helps counteract that blind auditions only achieve partial separation. Furthermore, it helps avoiding market failure, which may arise if neither informed nor blind auditions are profitable. In this case, it would seem natural for the evaluator to revert to the outside option immediately without holding an audition at all. This would lead to continual subcontracting or zero-hour contracts in the performing arts, reinforcing gender inequality via the difficulty of combining precarious employment, often characterised by "bulimic patterns of working", with motherhood and caring responsibilities (Conor et al., 2015).

Market failure may have adverse and wide-ranging consequences for welfare beyond gender inequality. One mechanism could be reduced civic engagement: research by the National Endowment For The Arts (2009) finds that"Americans who create or perform art are more civically active than the general U.S. adult population" (p.6). A second mechanism could be reduced spillover effects accruing to the neighbourhood in which musicians reside and work: research by the Social Impact of the Arts Project (2017) at the University of Pennsylvania, School of Social Policy and Practice, finds that "the presence of cultural resources in a neigho[u]rhood has a significant positive impact on a neighbo[u]rhood's health, the outcomes of its schools and its crime rate" (p.3). While the research group stresses that this is not a causal relationship and that the positive impact is more pronounced in moderate- and low-income neighbourhoods, they argue that cultural resources, such as artists, "are integral to a neighbourhood ecology that promotes social wellbeing" (p.3).

7

Avenues for Future Research

To conclude, I discuss avenues for future research which are beyond the scope of the thesis. In a first step, I plan to study the robustness of my model's results by relaxing various assumptions; for example, by allowing for more general assumptions on effort choices, distributions, and production functions. A more general effort choice captures the idea that applicants commonly prepare brief excerpts from a repertoire that the evaluator sets in advance. By introducing a general distribution in the evaluator's outside option, I intend to show that my results do not hinge on the uniform distribution assumption. Finally, I intend to demonstrate the robustness of my results under production functions that do not imply perfect substitutability between ability and effort beyond the generalisation analysed in Appendix A. One possibility for greater generality would be a constant elasticity of substitution production function.

In a second step, I plan to study the potential trade-off between policies mitigating biases and effort incentives for applicants in more depth. One example is a setting in which the evaluator's bias is not common knowledge. Some important research questions that would arise in such a setting are: what if applicants are uncertain or underestimate the evaluator's bias? What if the evaluator is partially sophisticated or naive about his bias when choosing the form of audition? This is motivated by evidence that many forms of bias work at a subconscious level. One example is the vision heuristic (Tsay, 2013): evaluators consistently report to value sound as central in performance but yet arrive at different winners depending on whether they have visual information about applicants. This suggests that "[p]rofessional judgment appears to be made with little conscious awareness that visual cues factor so heavily in preferences and decisions" (p.14583). Therefore, it is important to explore the implications if evaluator's bias is not known with certainty.

A further avenue is to examine welfare implications to obtain recommendations for policy. A starting point may be what a social planner with a utilitarian welfare function would do: under which relative weights on the types of applicants in social welfare does the planner prefer a blind audition? A second point of departure may be affirmative action if participation of females is lower ex-ante. In this case, can policies ensuring gender equality compensate for biases?

I intend to analyse the implications of introducing competition among applicants. A first point of departure may be a tournament model à la Cornell and Welch (1996). Another promising avenue for research is to endogenise the evaluator's bias. Rather than assuming a reduced-form bias parameter, a categorical model of cognition (Fryer and Jackson, 2008) may be a plausible way to microfound biased decision making on part of the evaluator. This would allow me to show that a bias against female applicants can emerge from the evaluator's sorting of past experiences into categories if such categorisation affects his decision making. This would also emphasise that a bias against female applicants can obtain even when the evaluator has no malevolent taste for discrimination.

Lastly, a replication study of Goldin and Rouse (2000) using UK data, or testing the hypothesis that the authors' findings are driven by taste-based rather than statistical discrimination, would underline the importance and empirical relevance of my model.

Appendix A

Multiplicative Technology

Suppose the production function, $f(q, \eta) = q\eta$, is multiplicative instead of additively separable. Analogously to the discussion in the main text, the random outside option is, thus, distributed $\bar{U} \sim U[1-\beta, 4]$.

First, note that this assumption implies supermodularity between effort and ability, and makes being of high ability more valuable to the evaluator. As a result, exerting high rather than low effort increases the probability of being hired by *less* for low-ability applicants in an informed audition compared to the additive technology. Exerting high rather than low effort increases the probability of being hired by *more* for high-ability applicants. Take, for example, $\beta = 1$. By exerting high rather than low effort in an informed audition, low-ability females increase their probability of being hired by $\frac{1}{3}$ under the additive, but only by $\frac{1}{4}$ under the multiplicative technology. In contrast, high-ability females increase their probability of being hired by $\frac{1}{3}$ under the additive, but by $\frac{1}{2}$ under the multiplicative technology.

In an informed audition, with the greater emphasis on ability, the threshold that high-ability applicants do not find it worthwhile to exert high effort increases from $\frac{6}{5}$ to $\frac{11}{7}$. Simultaneously, the threshold that low-ability males find it worthwhile to participate and exert low effort increases from $\frac{5-\sqrt{17}}{2}$ to 1. This implies that the interval of low biases in which only high-ability applicants participate in an informed audition, $\beta \in \beta_L^I \equiv [0, 1)$, is larger. Conversely, both the interval in which low-ability males join in and exert low effort, $\beta \in \beta_M^I \equiv [1, \frac{11}{7}]$, and the interval in which all applicants except low-ability females exert low effort, $\beta \in \beta_H^I \equiv (\frac{11}{7}, 2]$, is smaller.

In a blind audition, the greater emphasis on ability manifests in a significantly *reduced* range of priors for which a pooling equilibrium can be supported; that is, $\frac{7}{10} \leq p < 1$ with the upper bound on the bias given by $\beta(p) \equiv \sqrt{4p^2 + 30p - \frac{31}{4}} - 2p - \frac{5}{2}$. Intuitively, because ability is not revealed in a pooling blind audition, when ability is more important, the probability of being hired at the information set q_H is smaller. The evaluator is more cautious to hire an applicant when not being able to infer ability as it poses a greater risk to his revenues and he always has the option to break even with the outside option. As a result, for a given prior, it pays off less for low-ability applicants to pool with their high-ability counterparts. On the other hand, a separating equilibrium exists for a much *larger* range of biases, $\beta \in \beta_H^B \equiv (\frac{\sqrt{105}-9}{2}, 2]$, because the threshold that deters low-ability applicants from participating is significantly lower. Furthermore, because the evaluator can infer ability in a separating equilibrium and values ability more than under the additive technology, he finds it optimal to hire more often at q_H. It makes a separating blind audition more attractive to the evaluator.

When comparing the evaluator's net utility in both forms of audition for different values of the bias and prior, Proposition 4 modifies to:

Proposition 4′. *(i) For* $\frac{7}{10} \leq p < 1$*, if the evaluator's bias against female applicants is very low,* $\beta \in \beta_L^B \equiv [0, \beta(p)]$*, he prefers an informed audition over a pooling blind audition. (ii) For* $0 < p < 1$*, if the evaluator's bias against female applicants is low,* $\beta \in \beta_L^I \cap \beta_H^B = (\frac{\sqrt{105}-9}{2}, 1)$*, he prefers an informed audition over a separating blind audition. (iii) For* $0 < p < 1$*, if the evaluator's bias against female applicants is high,* $\beta \in \beta_H^I \equiv (\frac{11}{7}, 2]$*, he prefers a separating blind audition.*

The key difference compared to the additive technology in the main text is case (*ii*) in Proposition 4′; that is, the evaluator continues to be better off in an informed audition relative to a separating blind audition, as long as his bias is low enough as to not induce low-ability males to participate. While both forms of audition only attract high-ability applicants for this range of bias, the evaluator benefits from hiring high-ability males more often than their female counterparts. However, the evaluator has no means of distinguishing high-ability males from females in a blind audition; they are essentially the same type to him behind a curtain. As a result, he has to hire high-ability males and females with the same probability which necessarily makes him worse off. Taken together, for a bias below some threshold, this variant of the model predicts that the evaluator always prefers an informed audition. This is illustrated in the three panels of Figure A.1 for different priors p and threshold $\beta_L^I < 1$.

When comparing the applicant's preferences over auditions, the results for females and high-ability males remain qualitatively unchanged. For low-ability males, Proposition 7 modifies to:

Proposition 7′. *(i) A low-ability male prefers a blind audition if the evaluator's bias against female applicants is very low; that is* $\beta \in \beta_L^B \equiv [0, \beta(p)]$*. (ii) He is indifferent if the evaluator's bias against female applicants is low; that is* $\beta \in \beta_L^I \cap \beta_H^B = (\frac{\sqrt{105}-9}{2}, 1)$*. (iii) He*

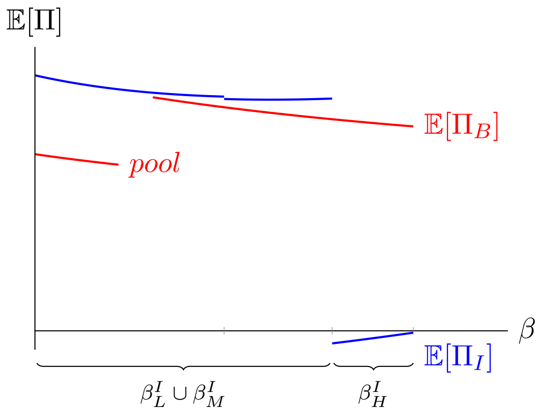

(a) $\Pr(\eta_H) = 0.9$

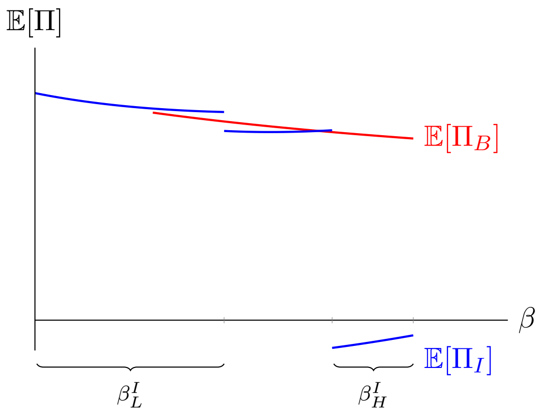

(b) $\Pr(\eta_H) = 0.5$

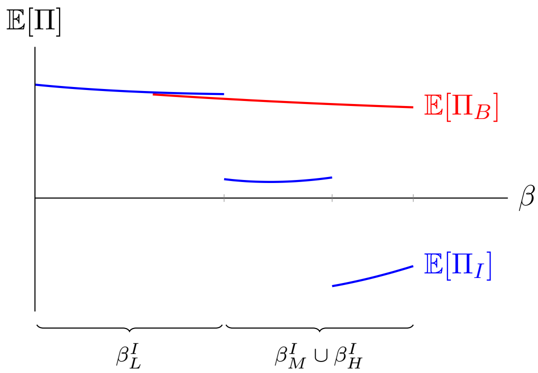

(c) $\Pr(\eta_H) = 0.1$

Figure A.1: Evaluator's Expected Net Utility in Blind and Informed Audition

prefers an informed audition if the evaluator's bias against female applicants is high; that is, $\beta \in \beta_M^I \cap \beta_H^B = [1, 2]$.

For a low bias, the intuition behind Proposition 7′ is akin to the one in the main text. The key difference compared to the additive technology is case (*ii*) because the threshold of β that induces a low-ability male to participate is much greater. As a result, beyond

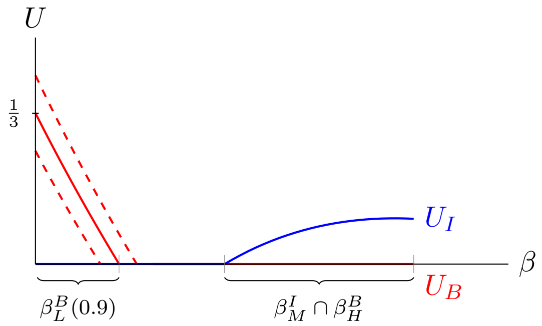

Figure A.2: Expected Utility of Low-ability Male in Blind and Informed Audition

a critical value, the bias is too strong to pool in a blind audition due to high effort costs. Yet, in an informed audition, the bias is not strong enough to make low effort worthwhile. Therefore, for this range of biases, he does not have an incentive to participate in either form of audition. Above this threshold, his preferences reverse as under the additive technology because he can reap the benefits from the evaluator's bias against female applicants in an informed audition.

Appendix B

Refinements in Separating Blind Audition

For the evaluator's belief that only high-ability applicants exert high effort to be consistent, he needs to have a sufficiently high bias; that is, $\beta \in \beta_H^B \equiv (\sqrt{17}-3, 2)$. In this case, low-ability applicants do not participate and the information set q_L is not reached. Therefore, the evaluator's belief at q_L is not determined by equilibrium play; in other words, perfect Bayesian equilibrium does not place restrictions on this belief as it is off the equilibrium path.

Furthermore, note that the Intuitive Criterion (Cho and Kreps, 1987) is moot. In the separating blind audition, exerting low effort is equilibrium dominated for all types. In particular, for a low-ability applicant, the equilibrium payoff from dropping out, O, is greater than the highest possible payoff from exerting low effort. For a high-ability applicant, the equilibrium payoff from exerting high effort is greater than the highest possible payoff from exerting low effort. Therefore, the requirement that the evaluator's belief at q_L places zero probability on nodes that are reached only if an applicant plays an equilibrium dominated strategy does not apply. It is not possible for the belief at q_L to place zero probability on all four nodes simultaneously.

B.1 Refinement based on Deviation Payoffs

I argue, however, that the evaluator's belief that the applicant is of low ability when observing a deviation to q_L is a reasonable restriction. Suppose to the contrary that, for a given bias $\beta \in \beta_H^B$, in the separating blind audition, the evaluator also believes the applicant to be of high ability if she delivers a low-quality performance behind the curtain. This essentially means that, if a low-quality performance is to be observed, it is because a high-ability applicant has trembled and chosen an easy piece by accident, or she rests on her laurels of being believed to be of high ability irrespective of the difficulty of the piece performed. In other words, a low-quality performance is not observed because a low-ability applicant is attracted to the audition to exploit the evaluator's belief at q_L. This perturbation of the evaluator's belief implies that, he expects a higher gross utility:

$$\mathbb{E}[V|q_L] = q_L + \mathbb{E}[\eta|q_L] - \frac{\beta}{2} = 1 + 2 - \frac{\beta}{2}$$

from hiring the applicant at q_L. Therefore, the probability of being hired at this information set changes to:

$$\Pr(\text{hired}|q_L) = \Pr(\bar{U} \leq 3 - \frac{\beta}{2}) = \frac{2+\beta}{4+2\beta}$$

Note that the hiring probability at q_L is higher than the one under the evaluator's original belief, $\frac{\beta}{4+2\beta}$. Intuitively, if the evaluator believes the applicant behind the curtain to be of high ability when hearing an easy piece, he should be more willing to hire her and overlook her mistake. Due to the altered hiring probability, the applicant's effort incentives change to:

$$\begin{aligned} U(e_L, \eta_L) &= \frac{2+\beta}{4+2\beta} w - \frac{1}{2} \\ U(e_H, \eta_L) &= \frac{4+\beta}{4+2\beta} w - 2 \\ U(e_L, \eta_H) &= \frac{2+\beta}{4+2\beta} w - \frac{1}{4} \\ U(e_H, \eta_H) &= \frac{4+\beta}{4+2\beta} w - 1 \end{aligned}$$

where $w \equiv \mathbb{E}[\bar{U}] = 3 - \frac{\beta}{2}$.

First, consider a low-ability applicant under the perturbed belief at q_L. Before, this type was deterred from participating for $\beta \in \beta_H^B$. Now, exerting low effort is worthwhile for all β_H^B (see Figure B.1). The low-ability applicant is, in fact, exploiting the evaluator's belief; it plays out in her favour. With a slight abuse of notation, the utility gain from deviating from O to e_L is given by:

$$\big[U(e_L, \eta_L) - U(O, \eta_L)\big]\Big|_{\beta \in \beta_H^B} = 1 - \frac{\beta_H^B}{4}$$

which is monotonically decreasing in the bias for all possible values. Therefore, the lowest possible utility gain from deviating to e_L occurs as β approaches its maximum:

$$\lim_{\beta \uparrow 2} \big[U(e_L, \eta_L) - U(O, \eta_L)\big] = \frac{1}{2} \tag{B.1}$$

Second, consider a high-ability applicant under the perturbed belief at q_L. Before, this type was exerting high effort for $\beta \in \beta_H^B$. Now, exerting high effort is only worthwhile if the

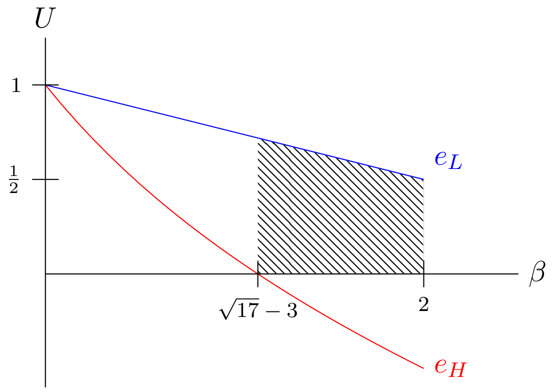

Figure B.1: Effort Choice of Low-Ability Applicant in Separating Blind Audition

bias does not exceed $\frac{6}{5}$. Exerting low effort is optimal for any bias exceeding this threshold (see Figure B.2). The high-ability applicant essentially avoids the additional effort costs and rests on her laurels by preparing an easy piece. With a slight abuse of notation, the utility gain from deviating from e_H to e_L is given by:

$$\left[U(e_L, \eta_H) - U(e_H, \eta_H)\right]\Big|_{\beta \in \beta_H^B} = \frac{5}{4} - \frac{4}{2 + \beta_H^B}$$

which is monotonically increasing in the bias for all possible values. Therefore, the greatest possible utility gain from deviating to e_L occurs as β approaches its maximum:

$$\lim_{\beta \uparrow 2} \left[U(e_L, \eta_H) - U(e_H, \eta_H)\right] = \frac{1}{4} \tag{B.2}$$

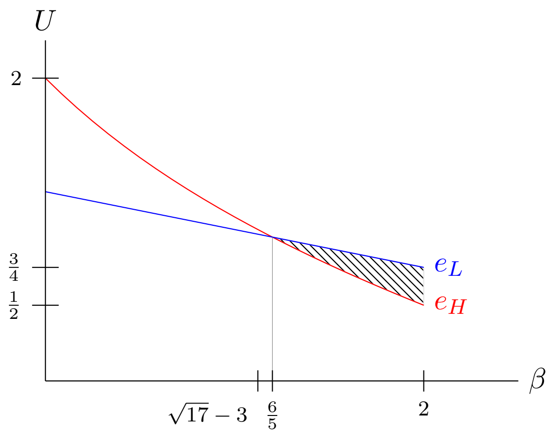

Figure B.2: Effort Choice of High-Ability Applicant in Separating Blind Audition

Finally, observe that for $\beta \in (\frac{6}{5}, 2)$, both high- and low-ability applicants gain from deviating to low effort under the perturbed belief. The utility gain for low-ability from deviating, however, is always greater than the utility gain for high-ability applicants from deviating. This follows immediately because the lowest possible utility gain for the former (B.1) exceeds the largest possible utility gain for the latter (B.2). As a result, the evaluator should say to himself that the applicant behind the curtain is more likely to be a low-ability applicant exploiting his belief. Put differently, a high-ability applicant trembling to an easy piece or resting on her laurels of being believed to be of high ability irrespective of the difficulty of the piece is less likely.

Observe also that for $\beta \in (\sqrt{17} - 3, \frac{6}{5}]$, only low-ability applicants have an incentive to deviate to low effort under the perturbed belief. High-ability applicants exert high effort. For this range of bias, by the Intuitive Criterion, the evaluator should believe the applicant to be of low-ability with probability one when observing a low-quality performance.

B.2 Refinement based on D1-Criterion

Under D1, similar to the Intuitive Criterion, deviations of the applicant emerge as the consequence of a rational decision. D1, however, is capable to put more structure on the evaluator's belief at q_L in the separating blind audition; it restricts the off-path belief to be a point belief with all mass on the type who is most likely to deviate.

Define $D(\theta, T, e_L)$ to be the set of the evaluator's mixed strategy best responses to action e_L and beliefs concentrated on T that make type θ strictly prefer e_L to her equilibrium strategy. Let $D^0(\theta, T, e_L)$ be the set of mixed best responses that make type θ exactly indifferent.

Definition 1. A type θ is deleted for strategy e_L under the D1-Criterion if there is a θ' such that

$$\{D(\theta, \Theta, e_L) \cup D^0(\theta, \Theta, e_L)\} \subset D(\theta', \Theta, e_L).$$

From Definition 1 follows that, if the set of the evaluator's mixed best responses (that is, the range of hiring probabilities) that make high-ability applicants willing to deviate to low effort is strictly smaller than the set of best responses that make low-ability applicants willing to deviate, then the evaluator should believe that low-ability applicants are infinitely more likely to deviate to low effort than their high-ability counterparts are. In other words, D1 tests a deviation to low effort for a given type with respect to each particular mixed best response of the evaluator (Fudenberg and Tirole, 1991, p.451-452).

Proposition 12. *After observing q_L, there are more best responses of the evaluator that improve the expected equilibrium utility of low-ability compared to high-ability applicants. Therefore, the evaluator should infer that he deals with the former and put zero weight on the latter.*

Proof. First, consider low-ability males and females in the separating blind audition. The minimal hiring probability at the low-performance information set, that induces these types to deviate from O to e_L, is given by:

$$\Pr(\text{hired}|q_L)w - \frac{1}{2} \geq 0$$
$$\Rightarrow \Pr(\text{hired}|q_L) \geq \frac{1}{6-\beta} \tag{B.3}$$

Second, consider high-ability males and females in the separating blind audition. The minimal hiring probability at the low-performance information set, that induces these types to deviate from e_H to e_L, is given by:

$$\Pr(\text{hired}|q_L)w - \frac{1}{4} \geq \Pr(\text{hired}|q_H)w - 1$$
$$\Rightarrow \Pr(\text{hired}|q_L) \geq \frac{4+\beta}{4+2\beta} - \frac{3}{12-2\beta} \tag{B.4}$$

where the probability of being hired at the high-performance information set is fixed at its equilibrium value, $\frac{4+\beta}{4+2\beta}$.

Finally, for a given bias, compare inequality (B.3) with inequality (B.4) to conclude that:

$$\{D((\eta_H,\cdot),\Theta,e_L) \cup D^0((\eta_H,\cdot),\Theta,e_L)\} \subset D((\eta_L,\cdot),\Theta,e_L)$$

Figure B.3 shows graphically that the set of the evaluator's mixed best responses to e_L that make low-ability applicants prefer low effort to their equilibrium strategy (light and dark grey area) strictly contains the set of mixed best responses for their high-ability counterparts (dark grey area) for all possible values of bias. By Definition 1, therefore, a deviation to low effort must be interpreted as coming from low-ability applicants.

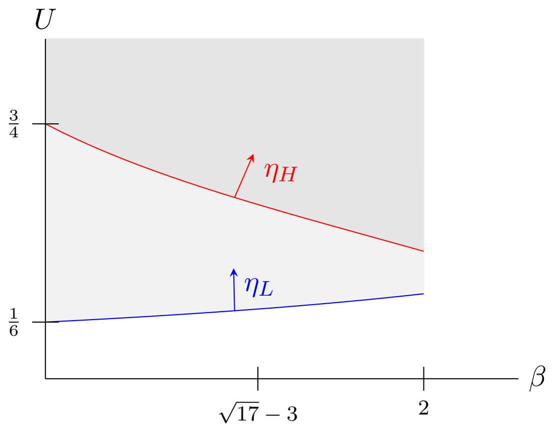

Figure B.3: Set of Evaluator's Mixed Best Responses to Low Effort inducing Deviation for Low- and High-Ability

□

Bibliography

Abeles, Hal, "Are Musical Instrument Gender Associations Changing?," *Journal of Research in Music Education*, 2009, *57* (2), 127–139.

Adler, Moshe, "Stardom and Talent," *The American Economic Review*, March 1985, *75* (1), 208–212.

Becker, Gary S., *The Economics of Discrimination*, 2 ed., Chicago: University of Chicago Press, 1971.

Bertrand, Marianne and Sendhil Mullainathan, "Are Emily and Greg More Employable Than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination," *The American Economic Review*, September 2004, *94* (4), 991–1013.

Cain Miller, Claire, "Is Blind Hiring the Best Hiring?," *The New York Times Magazine*, 25 February 2016. Available at https://www.nytimes.com/2016/02/28/magazine/is-blind-hiring-the-best-hiring.html. Accessed: 24 February 2019.

Castiglione, Lawrence V., "Performance, Practice, and Policy Considerations," *Design For Arts in Education*, September / October 1985, *1* (87), 31–37.

Chadwick, Evelyn, "Of Music and Men," *The Strad*, December 1997, pp. 1324–1329.

Cho, In-Koo and David M. Kreps, "Signaling Games and Stable Equilibria," *The Quarterly Journal of Economics*, May 1987, *102* (2), 179–221.

Clauhs, Matthew S., "The Effects of Race and Gender Bias on Style Identification and Music Evaluation." PhD dissertation, Temple University May 2013.

Conor, Bridget, Rosalind Gill, and Stephanie Taylor, "Gender and Creative Labour," *The Sociological Review*, 2015, *63* (S1), 1–22.

Cornell, Bradford and Ivo Welch, "Culture, Information, and Screening Discrimination," *Journal of Political Economy*, 1996, *104* (3), 542–71.

Elliott, Charles A., "Race and Gender as Factors in Judgments of Musical Performance," *Bulletin of the Council for Research in Music Education*, 1995, *127* (Winter), 50–56.

Ericsson, K. A., "The Influence of Experience and Deliberate Practice on the Development of Superior Expert Performance," in K. A. Ericsson, N. Charness, P. J. Feltovich, and R. R. Hoffman, eds., *The Cambridge Handbook of Expertise and Expert Performance*, Cambridge: Cambridge University Press, 2006, chapter 38, pp. 683–703.

_ , R. T. Krampe, and C. Tesch-Romer, "The Role of Deliberate Practice in the Acquisition of Expert Performance," *Psychological Review*, July 1993, *3* (100), 363–406.

Fryer, Roland and Matthew Jackson, "A Categorical Model of Cognition and Biased Decision Making," *The B.E. Journal of Theoretical Economics*, 2008, *8* (1), 1–44.

Fudenberg, Drew and Jean Tirole, *Game Theory*, Cambridge, MA: MIT Press, 1991.

Gallwey, W. Timothy, *The Inner Game of Tennis*, New York: Random House, 1974.

Gladwell, Malcolm, *Blink: The Power of Thinking without Thinking*, New York: Little, Brown & Company, 2005.

Goldin, Claudia and Cecilia Rouse, "Orchestrating Impartiality: The Impact of "Blind" Auditions on Female Musicians," *The American Economic Review*, September 2000, *90* (4), 715–741.

Grossman, Sanford J., "An Introduction to the Theory of Rational Expectations Under Asymmetric Information," *Review of Economic Studies*, October 1981, *48* (4), 541–559.

Hart, Philip, *Orpheus in the New World: The Symphony Orchestras as an American Cultural Institution its Past, Present, and Future*, New York: W.W. Norton & Company, 1973.

Hays, Kate F. and Charles H. Brown, *You're on! Consulting for Peak Performance*, Washington, DC: American Psychological Association, 2006.

Holland, Bernard, "The Fair, New World of Orchestra Auditions," *The New York Times*, 11 January 1981.

Kamakura, Wagner A. and Carl W. Schimmel, "Uncovering Audience Preferences for Concert Features from Single-Ticket Sales with a Factor-Analytic Random-Coefficients Model," *International Journal of Research in Marketing*, June 2013, *30* (2), 129–142.

Laffont, Jean-Jacques and Jean Tirole, "Using Cost Observation to Regulate Firms," *Journal of Political Economy*, 1986, *94* (3), 614–641.

Levinson, Gary, "Preparing for an Orchestral Audition: the Basics," https://www.thestrad.com/preparing-for-an-orchestral-audition-the-basics/6964.article July 2017. Accessed: 22 February 2019.

McPherson, Gary E. and Emery Schubert, "Measuring Performance Enhancement in Music," in Aaron Williamon, ed., *Musical Excellence: Strategies and Techniques to Enhance Performance*, Oxford: Oxford University Press, 2004, chapter 4, pp. 61–82.

Meyer, Margaret A. and John Vickers, "Performance Comparisons and Dynamic Incentives," *Journal of Political Economy*, 1997, *105* (3), 547–581.

Milgrom, Paul R., "Good News and Bad News: Representation Theorems and Applications," *The Bell Journal of Economics*, Autumn 1981, *12* (2), 380–391.

Musicians' Union, "Employment Agreements: Negotiated Rates," https://www.musiciansunion.org.uk/Home/Advice/Playing-Live/Orchestral-Work/Pay-Conditions/Employment-Agreements. Accessed: 23 February 2019.

National Endowment For The Arts, "Art-Goers in Their Communities: Patterns of Civic and Social Engagement," https://www.arts.gov/sites/default/files/98.pdf October 2009. Accessed: 09 March 2019.

Rosen, Sherwin, "The Economics of Superstars," *The American Economic Review*, December 1981, *71* (5), 845–58.

Seltzer, George, *Music Matters: The Performer and the American Federation of Musicians*, Mettuchen, NJ: Scarecrow Press, 1989.

Social Impact of the Arts Project, "Culture and Social Wellbeing in New York City: Highlights of a Two-Year Research Project," https://repository.upenn.edu/siap_culture_nyc/2 February 2017. Accessed: 09 March 2019.

Spence, Michael, "Job Market Signaling," *The Quarterly Journal of Economics*, August 1973, *87* (3), 355–374.

Starr, Susan, "The Prejudice Against Women," *Music Journal*, 1 March 1974, *32* (3), 14–15, 28.

Taylor, Curtis R. and Huseyin Yildirim, "Subjective Performance and the Value of Blind Evaluation," *The Review of Economic Studies*, April 2011, *78* (2), 762–794.

The Economist, "All Ears," *The Economist*, 30 November 1996, *341* (7994), 89–90; UK 135–136.

Thorngate, Warren, Robyn M. Dawes, and Margaret Foddy, *Judging Merit*, New York: Psychology Press, 2010.

Towse, Ruth, "Market Value and Artists' Earnings," in Arjo Klamer, ed., *The Value of Culture: On the Relationship between Economics and Arts*, Amsterdam: Amsterdam University Press, 1996, chapter 6, pp. 96–107.

_ , *A Textbook of Cultural Economics*, Cambridge, UK: Cambridge University Press, 2010.

Tsay, Chia-Jung, "Sight over Sound in the Judgment of Music Performance," *PNAS*, 3 September 2013, *110* (36), 14580–14585.

_ **and Mahzarin R. Banaji**, "Naturals and Strivers: Preferences and Beliefs about Sources of Achievement," *Journal of Experimental Social Psychology*, 2011, *47* (2), 460–465.

Universal Music Group, "Universal Music Group signs Global Recording, Artist Services Agreement for New TV Music Sensation, 'The Voice of ...'," https://www.universalmusic.com/universal-music-group-signs-global-recording-artist-services-agreement-for-new-tv-music-sensation-the-voice-of/ 19 January 2011. Accessed: 22 February 2019.

Vernos, Isabelle, "Quotas are Questionable," *Nature*, 7 March 2013, *495*, 39.