

On the magnification of small biases in decision-making¹

Shaun William Davies²
Edward Dickersin Van Wesep³
Brian Waters⁴

September 20, 2019

Abstract

We analyze a setting in which an actor chooses between N *ex ante* identical options. She can exert effort to learn about the quality of each option, but can ultimately choose only one. There are up to $N!$ unique optimal effort vectors, and each is asymmetric: a large amount of effort is expended learning about one arbitrarily chosen option, less on another, even less on a third, etc. This implies asymmetric likelihoods of each item being chosen. If the actor has an infinitesimal bias in favor of one option, then the actor selects an effort vector that maximizes the likelihood of her favored option being chosen. Small biases are magnified, sometimes enormously. We also show that a glass ceiling can arise, in which favored types are increasingly prevalent as one ascends the corporate ladder. These results have implications for portfolio selection (e.g., home bias, socially responsible investment funds), hiring (e.g., CEO choice, the glass ceiling), start-up funding, and a variety of other applications.

Keywords: choice, bias, portfolio selection, glass ceiling

JEL: D83; D91; J71; G41

¹We would like to thank conference and seminar participants at the University of Amsterdam, the University of Colorado, Boulder, the 20th Meeting of the Finance Theory Group, and the UCLA Anderson Finance Conference (In Honor of Bhagwan Chowdhry) for their helpful insights and suggestions. All errors are our own.

²Leeds School of Business, University of Colorado, Boulder, Campus Box 419, Boulder, CO 80309, shaun.davies@colorado.edu.

³Leeds School of Business, University of Colorado, Boulder, Campus Box 419, Boulder, CO 80309, Edward.VanWesep@colorado.edu.

⁴Leeds School of Business, University of Colorado, Boulder, Campus Box 419, Boulder, CO 80309, brian.waters@colorado.edu.

1 Introduction

Many choice settings require an actor to select a single item from a set of options. Examples include a board choosing a new chief executive officer (CEO), a firm hiring an employee, an activist investor choosing a firm to target, a finance department choosing a new assistant professor, an entrepreneur choosing an idea for a new firm, a landlord choosing a tenant, and a consumer choosing an insurer. These settings differ from those in which any subset of N options can be chosen because choosing one inherently requires not choosing others. Any effort toward learning about one option is therefore either a complement or a substitute for efforts toward learning about other options.

Decision-makers in these settings are often accused of bias, and many of these accusations have some bite. For example, most S&P 500 CEOs are male, and the fraction of males rises as one moves up the corporate ladder. This “glass ceiling” is often attributed to bias in how people view female leaders. Credible claims of bias in hiring extend well beyond the C-suite. For example, female musicians are more likely to be hired when auditions are blind (Golden and Rouse, 2000) and job applicants with stereotypically black names are less likely to receive callbacks (Bertrand and Mullinathan, 2004).¹ Anti-nepotism laws, designed to mitigate or eliminate bias in hiring, are prevalent. The perception of bias is not unique to hiring, however, and may be a product of the choice setting. Landlords, for example, are frequently accused of bias in their choices of tenants (Fred and Donald Trump being notable examples).² Angel and venture funders are also accused of bias in funding male-led ventures.³

Is there a connection between perceived bias and the specifics of the choice problem?⁴ In this paper, we show that the optimal behavior of an actor facing this choice problem magnifies any infinitesimal biases that may exist. It does not take much bias to generate large disparities in outcomes.

We develop a simple model in which an actor must select one item from N options, and her

¹See Blau and Kahn, 2017, for a recent review of the gender pay gap literature.

²<https://www.nytimes.com/2016/08/28/us/politics/donald-trump-housing-race.html>

³John Doerr, a partner at the Venture Capital Firm Kleiner Perkins famously said “That correlates more with any other success factor that I’ve seen in the world’s greatest entrepreneurs. If you look at Bezos, or Andreessen, David Filo, the founders of Google, they all seem to be white, male, nerds who’ve dropped out of Harvard or Stanford...” This statement was used as evidence of gender bias in a wrongful termination lawsuit against the firm.

⁴Perhaps the most obvious link is that, when only one item is selected, there is no *statistical* way to demonstrate bias: it is always possible that the best choice happens to be of a particular type. Perhaps, for example, one can argue that there is bias in favor of male CEOs, but there is no way to present statistical evidence that a specific board was biased in appointing a specific male CEO.

prior regarding the quality of each item is identical. She simultaneously chooses a set of efforts that generate signals of each item’s quality. After observing these signals, she forms posteriors and chooses a single item based on those posteriors. We show, perhaps surprisingly, that an optimal effort vector in the simultaneous choice problem is asymmetric: the actor arbitrarily chooses one item and invests substantial effort toward learning the quality of that item; she arbitrarily chooses a second item and invests less effort toward that item, even less toward a third, etc. Because the assignment of efforts to items is arbitrary, there are up to $N!$ optimal effort vectors, each with the same set of efforts but a different assignment of those efforts to items.⁵

Asymmetric effort generates differential likelihoods of each item being chosen. While the actor’s expected payoff is identical for each of the $N!$ optimal effort vectors, the likelihood of a particular item being chosen is not identical. For one vector, the likelihood of item 6 being chosen might be 20%, and for another vector only 1%. From the “item’s” perspective, the choice of effort vector matters.

Next, we allow the actor to have an infinitesimal bias in favor of one particular item.⁶ This bias eliminates the arbitrariness of her effort assignment. A subset of assignments is optimal and these are the ones that maximize the likelihood of her favored item being chosen. As a result, her favored item is selected more often than $1/N$ of the time, often much more. Depending on the information structure, her favored item may be chosen more than *half* of the time, even if N is large! We therefore show that small biases can be magnified when the actor must choose only one option from a set. Furthermore, this implies that large differences in the rates with which certain items are selected do not imply large biases.

This result has two important implications. The first implication is relevant for policy-makers. Simply reducing the size of the bias will not necessarily have much effect on the magnitude of the “selection gap.” If policy-makers want to reduce the gap, then, for example, educating decision-makers about implicit bias may not be very useful. The second implication is important for researchers. Suppose that researchers find persistent gaps in selection likelihoods but only observe small biases

⁵We model the problem as one of simultaneous choice. An alternate model would require sequential choice, as in the literature on multi-armed bandits. We believe that the main intuition of the model would follow in that case as well. Indeed, it’s more surprising in the simultaneous case because all but the first action in a bandit problem can be history-dependent.

⁶We model the bias as a private benefit that accrues to the actor if she chooses her preferred option. This is equivalent to assuming that her prior distribution for her preferred option is equal to that of the other distributions in all moments except for a higher mean.

among selectors. This fact does not establish that bias is not generating the selection gaps.

Finally, we consider a highly simplified version of the model that permits closed-form solutions for the optimal effort vector and we show that the magnification of the bias can be higher when the importance of choosing a high quality item is higher. The reason is that higher stakes mean across-the-board higher efforts, which can increase the odds of the favored item being chosen. We call this result the “glass ceiling” effect: if we assume that as one ascends the corporate ladder that the importance of employee quality increases, and if we assume a small bias on the part of hiring managers in favor of white men, then our results suggest that we should see more white men as we ascend the corporate ladder. Indeed, we do.

It may be surprising that allowing the actor to undertake effort to learn about the quality of each item does not overwhelm the bias – in fact, those efforts are what magnify the bias in the first place. It may also be surprising that, when the stakes are high, the bias is still important (Becker, 1971, argues against this possibility). We show that higher stakes can actually exacerbate the problem.

If one wishes to create models with counter-intuitive results, it is not particularly difficult so long as one has enough freedom with assumptions. This model, however, is *the simplest model that one could write concerning the decision to choose one item from N* . In its simplest form (which we analyze after developing the more general model), there are two types of item (high and low quality), two signals that can result from learning (high and low signals), and two items from which to choose. A simpler model could not be written with learning and choice. Yet, in this model, we see all of the results discussed above.

Taking our theoretical predictions to the data is beyond the scope of the paper, but we can evaluate others’ work through the lens of our model. Hebert (2019), for example, studies the likelihood of an entrepreneur receiving funding conditional on her or his gender and the distribution of genders in her or his chosen industry. She shows that women are more likely than men to receive investment if and only if their proposed venture is in an industry in which female-led start-ups are more common. She also finds that the performance of female-led start-ups is higher than the performance of male-led start-ups if and only if the associated industry is predominantly male. For example, a woman starting a software firm is less likely than a man to receive funding, but enjoys superior average *ex post* performance. A woman starting a clothing line is more likely than a man

to receive funding, but suffers inferior average *ex post* performance.

Our model is consistent with these empirical findings. Suppose that there is a bias in favor of men and suppose that most candidates for equity funding are not deserving. If the actor is deciding between one male and one female candidate, the actor will, due to her or his scepticism, exert more effort learning about a female candidate. And, as we show in the analysis, she or he will exert no effort learning about the male candidate. If the woman is chosen, then it means that her signal is positive and the posterior for the likelihood that she is high quality is above the prior. If the man is chosen, then it means that the woman's signal is negative – the posterior for the selected man's quality is the same as the prior and therefore lower than the posterior for the woman if she is chosen.

Our model also suggests that, if we were to examine the long-run success of *unfunded* women and men, unfunded women who attempted a venture in a male dominated industry would perform worse than unfunded men who attempted a venture in a male dominated industry. An unfunded woman is associated with a bad signal whereas an unfunded man is not. The posterior for her quality is lower than for his, both conditional on not being chosen.

This result is contrary to what one obtains with a model in which a lower threshold for perceived quality is applied to favored groups. That is, one could easily imagine that a member of a favored group will receive funding if her expected quality is above x while a member of a disfavored group will receive funding if the expected quality is above y , in which $x < y$. In that model, the *ex post* performance, conditional on funding or conditional on non-funding, is lower for the favored group. Which model is correct is an empirical question.

In Section 2, we discuss the related literature. In Section 3, we present the base model and show that the optimal effort vector is asymmetric, meaning that the likelihood of a given item being chosen is different for different items even though they are *ex ante* identical. We also show that, once the decision-maker has a small bias, she assigns her efforts to maximize the likelihood of the favored item being chosen. Small biases are magnified, possibly by many orders of magnitude. In Section 4, we restrict attention to the model's simplest form so that we can derive statements for the optimal effort vectors. This allows us to calculate the likelihood of each item being chosen and show how this varies with parameters. We show that increasing the importance of choosing a good item increases the magnification effect. Section 5 concludes. Where they do not appear in the text,

proofs can be found in Appendix A.

2 Literature

The literature that relates to this manuscript is extensive. For the sake of brevity, we point to classic papers in the field as well as to surveys.

Our choice setting is most closely related to the so-called “bandit” literature, in which an actor can experiment with different slot machines *ad infinitum* (Rothschild, 1974, was seminal. See Bergmann and Valimaki (2006) for a review). She is never obligated to choose a single machine, but typically one machine will pay off often enough that she will stick with that machine forever. In some sense, bandit problems are sequential learning versions of our model, which has simultaneous learning. At least two implications of bandit models parallel our findings. First, if the distributions of payoffs for each option are *ex ante* identical, then the likelihood that the actor chooses a particular option is not constant across options. Whatever “arms” the actor pulls initially lead to path dependence in the actor’s eventual choice. Second, the actor might not end up with the best possible choice.

For the applications that we discussed in the introduction, learning is most likely some combination of simultaneous and sequential learning. One does not usually choose a CEO after a five-year trial run, but nor does one choose a CEO after one round of interviews. That said, the similarity of some of the models’ implications with those from the bandit literature suggests that our results are robust.

Beyond the bandit literature, there is a theory literature concerning the interaction of small differences in initial beliefs, information, or bias, and the effect of learning on those differences. In the same vein of our manuscript, for example, van Nieuwerburgh and Veldkamp (2008) find that small informational differences concerning asset values are amplified by learning: it is profitable to focus efforts toward acquiring information that others do not have. van Nieuwerburgh and Veldkamp (2008) apply this intuition toward explaining home bias in portfolio choice, and extend the concept to under-diversification more broadly in van Nieuwerburgh and Veldkamp (2010).

A theory literature has also arisen on the presence of biased beliefs. If learning the truth is important, how and why do biased beliefs persist? For example, Bordalo, Coffman, Gennaioli, and

Shleifer (2016) model stereotypes as arising from an overweighting of representative types. If two populations are very similar, but differ in one or two areas, people tend to identify those populations with those areas (e.g., Floridians are elderly). In similar work, Gennaioli and Shleifer (2010) model beliefs that are biased by limited recall and allow for learning. Their model's results are consistent with several judgement biases as well as portfolio anomalies.

In work closely related to the model in this manuscript, Bartos, Bauer, Chytilova, and Matejka (2016) run field experiments in which decision makers can exert effort in reading applications. They provide a simple model in which there is an initial bias against one group that is magnified by endogenous effort toward learning. They find that decision makers spend less time reading resumes from applicants with minority sounding names but more time evaluating rental applications. These experimental results correspond to the results in our paper.

A substantial literature in psychology has documented the prevalence of conscious and subconscious biases in many areas. Biases may be harmful for society, but arise because they are helpful for individuals or "tribes". For example, Kahneman (2011) summarizes decades of his own work, arguing that biases are useful for quick thinking. For example, if one sees a large, dark shape moving behind a bush, the cost of an error in estimating the likelihood that the motion is due to a dangerous animal is not symmetric. If I believe that it is more likely to be dangerous then I may waste energy, but if I believe that it is less likely to be dangerous, I may be dead. Kahneman (2011) also provides an extensive review of the literature on biases.

Of particular interest to us, the gender gap, whether it is the fraction of women in a particular field or at a particular level in the corporate hierarchy, has been extensively studied, as has the associated pay gap. The pay gap appears to result mostly from differential educational attainment (a difference that has now disappeared in the United States), experience on the job, hours worked, occupational choice, negotiating approach, employer bias, and health insurance costs, though some gap remains unexplained.

The literature has not determined the relative importance of each, in part because many interact. An example that is particularly relevant for our paper is occupational choice: the fact that the fraction of women differs by occupation is likely to be, at least in part, a result of bias by employers, so attributing gaps to occupational choice rather than bias is misleading. Bertrand (2011) and Blau and Kahn (2017) provide partial summaries of this literature.

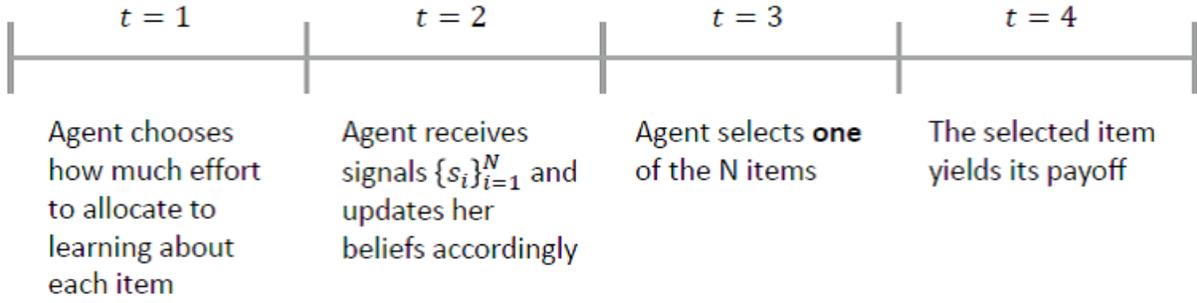


Figure 1: Timeline.

3 Base Model

An actor must select one of $N \geq 2$ options, indexed $i \in I \equiv \{A, B, \dots\}$, each of which has an unknown type $\theta_i = \theta \in \Theta \subset \mathbb{R}$, where we assume for simplicity that Θ is a countable set. While the actor does not initially know each option's type, she knows the prior probability, $p(\theta) = \Pr(\theta_i = \theta)$, that $\theta_i = \theta$, which is independent of i . In addition, for each option i , she can exert costly effort $e_i = e \in X \equiv [0, \bar{e}]$ to learn about the item's underlying type. After exerting efforts, for each option i , the actor observes a signal $s_i = \psi \in \Psi$, where Ψ is a random variable whose probability mass function (pmf), conditional on θ and e , is $q_i(\psi|\theta, e) = \Pr(s_i = \psi|\theta_i = \theta, e_i = e)$. We use the subscript i , since the probability $q_i(\cdot)$ depends on the realization of θ_i and on the effort exerted learning about item i .

Let $g_i(\psi|e) = \sum_{\Psi} p(\theta)q(\psi|\theta, e)$ denote the probability that $s_i = \psi$ condition on effort $e_i = e$, and let $r_i(\theta|\psi, e) = \frac{p(\theta)q_i(\psi|\theta, e)}{g_i(\psi|e)}$ denote the actor's posterior belief that $\theta_i = \theta$ after observing signal $s_i = \psi$ conditional on effort $e_i = e$. This implies that $E[\theta_i|s_i = \psi|e_i = e] = \sum_{\Theta} r_i(\theta|\psi, e)\theta$.

After observing these signals, the actor makes a choice $y \in I$, and payouts result. The timeline of the game is displayed in Figure 1.

The actor's payoff conditional on selecting item y is

$$U = \theta_y - \sum_{i \in I} c(e_i),$$

in which $c(e_i)$ is the cost of exerting effort to learn about item i with $c(e), c'(e), c''(e) > 0$ for $e > 0$ and $c(0) = c'(0) = 0$. Note that we do not assume direct complementarity or substitutability of

efforts in the cost function.

We make three additional assumptions regarding the relationship between effort e_i and the realized signal s_i . These assumptions are stronger than necessary but they allow for shorter and clearer derivations of our results.

Assumption 1: $q(\psi|\theta, e)$ is continuously differentiable with respect to effort e for all $e \in X$, $\psi \in \Psi$, and $\theta \in \Theta$.

Assumption 1 guarantees that the relationship between effort and the realized signal is well-behaved, and in particular guarantees that both $\frac{d}{de_i} \Pr(s_i = \psi|e_i)$ and $\frac{d}{de_i} E(\theta_i|s_i, e_i)$ exist and are continuous whenever,⁷

$$\Pr(s_i = \psi|e_i) = \sum_{\theta \in \Theta} \Pr(s_i = \psi|\theta_i = \theta, e_i = e)q_\theta > 0.$$

In plain English, we assume that, if signal ψ ever occurs for a given effort level, then both its likelihood and the expected payoff are continuously differentiable in effort.

Assumption 2: If $g_i(\psi|0) > 0$, then $E(\theta_i|s_i = \psi, e_i = 0) = E(\theta_i)$.

Assumption 2 implies that learning is never free. If the actor exerts no effort learning about a particular option, her posterior expected value after observing her signal must be equal to the prior expected value.

Assumption 3: For all $e \in X$, there exists $\psi \in \Psi$ such that $g_i(\psi|e) > 0$ and $\frac{d}{de_i} E(\theta_i|s_i = \psi, e_i) \neq 0$.

Assumption 3 ensures that at least one signal ψ (i) occurs with positive probability, and (ii) is differently informative for different levels of effort. In plain English, we assume that there exists some signal that occurs sometimes and is stronger or weaker for different levels of effort. Assumption 1 through 3 ensure that the effort problem is not degenerate and are sufficient to derive our key results in this section when there are only two options ($N = 2$). If there are more than two options, we require a slightly stronger assumption. Let $\underline{\psi}(e_i)$ denote the set of “lowest” possible signals for each level of effort $e_i \in X$,

$$\underline{\Psi}(e_i) = \arg \min_{\psi \in \Psi} E(\theta_i|s_i = \psi, e_i),$$

⁷The conditional expectation $E(\theta_i|s_i, e_i)$ is given by $E(\theta_i|s_i, e_i) = \sum_{\theta \in \Theta} \left(\frac{\Pr(s_i = \psi|\theta_i = \theta, e_i = e)q_\theta}{\Pr(s_i = \psi|e_i)} \right) \theta$.

and denote $\underline{\theta}$ to be the maximum conditional expected value that can be obtained across all levels of effort when realizing a lowest possible signal,

$$\underline{\theta} = \max_{e_i \in X} E(\theta_i | s_i = \psi, e_i) \text{ subject to } \psi \in \underline{\Psi}(e_i).$$

Assumption 3b: For all $e_i \in X$, there exists $\psi \in \underline{\Psi}$ such that $\Pr(s_i = \psi | e_i) > 0$ and $\frac{d}{de_i} E(\theta_i | s_i = \psi, e_i) \neq 0$ with $E(\theta_i | s_i = \psi, e_i) \geq \underline{\theta}$.

When there are greater than two options, the signal that satisfies Assumption 3 for each level of effort must impact the actor's expected payoff. If the $E(\theta_i | s_i = \psi, e_i) < \underline{\theta}$, then the actor may never choose option i when realizing signal $s_i = \psi$, and the signal will have no bearing on the actor's eventually payoff. Again, the assumption is to ensure that the effort problem is not degenerate for option i .

Figure 2 presents an example signal structure with four signals: $\Psi = \{\psi_1, \psi_2, \psi_3, \psi_4\}$. For all four, $\frac{d}{de_i} E(\theta_i | \psi_j, e_i) \neq 0$ and, because there are a finite number of signals, some subset must occur with positive probability. Therefore, the signal structure in Figure 2 satisfies our requirements.

Many standard information structures satisfy these assumptions, and we provide examples in the Appendix. An example to which we repeatedly turn in this paper is the simple case of two signals, $s_i \in \{L, H\}$, that satisfy

$$\Pr(s_i = H | \theta_i \geq \bar{\theta}) = \Pr(s_i = L | \theta_i < \bar{\theta}) = \frac{1 + e_i}{2},$$

and

$$\Pr(s_i = H | \theta_i < \bar{\theta}) = \Pr(s_i = L | \theta_i \geq \bar{\theta}) = \frac{1 - e_i}{2},$$

for some $\bar{\theta} \in \Theta$.

It is trivial to show that $\frac{d}{de_i} E(\theta_i | H, e_i) > 0$ and $\frac{d}{de_i} E(\theta_i | L, e_i) < 0$. It is also trivial to show that at least one of the two occurs with positive probability (and if $e_i < 1$ and $\bar{\theta}$ is within the support of θ_i , then both occur with positive probability). Therefore, this information structure also satisfies our assumptions. Finally, we assume that $c(e_i)$ is sufficiently convex so that the actor's problem is a well-behaved, concave optimization.

Before presenting our first results, we briefly discuss the preceding assumptions. First, while

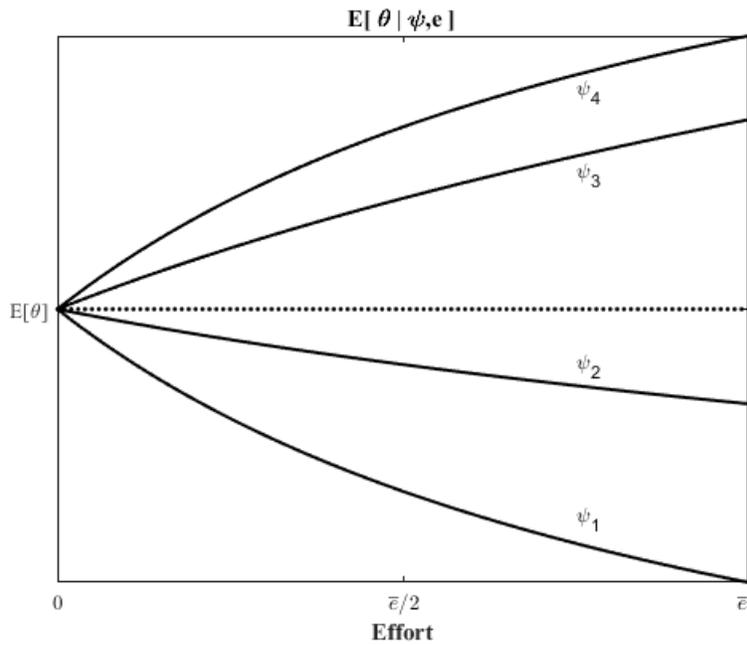


Figure 2: This example is of four signals, each of which is differentially informative for different levels of effort. Signals are ordered from best to worst, and the ordering does not depend on the level of effort. If effort is 0, then all signals are pure noise so there is no free learning. As effort increases, “good” signals get better and “bad” signals get worse, satisfying Assumption 3. Assumption 1, of continuous differentiability, appears to hold.

for simplicity we assume the state space is finite, it should be clear in what follows that the state space is essentially unrestricted. We could assume two types of item, 20 types, or a continuum of types. The set of possible types is not important for our results, so long as $E(\theta)$ exists and is finite. Second, efforts are not direct compliments or substitutes. This works to our disadvantage, as the key result of the model is that, when a single item must be chosen, efforts are *indirect* compliments or substitutes (depending on the information structure). Simply assuming that they are direct compliments or substitutes would generate the same results, but exogenously rather than endogenously. Third, the signal space must allow positive mass on at least one “relevant” signal. This assumption matters. As we show below, the results depend critically on the fact that, when efforts are the same, it is possible that two signals are also the same. It does not matter what is the likelihood of this occurring, but it does matter that it could occur.

3.1 Asymmetric efforts

The model is simple and stylized, but hopefully hews closely to the problems that decision-makers face in our applications. The first result follows immediately.

Proposition 1. *There are up to $N!$ unique optimal effort vectors, each of which yields the same expected payoff. Each vector is associated with a mapping $M : \{A, B, \dots\} \rightarrow \{1, 2, \dots, N\}$ in which $e_1 \geq e_2 \geq \dots \geq e_N$ (strict for $e_n < \bar{e}$).*

Proposition 1 establishes the form of the optimal effort vector. Focusing on the interior case in which no efforts are ever at their maximum allowed values, the actor chooses one item and devotes the most effort to that item. She chooses a second item and devotes less effort than to the first. She chooses a third and devotes less effort than to the second, etc., until she is left with a final item to which she devotes the least effort. There are $N!$ equally profitable effort vectors because the assignment of items to efforts is arbitrary. Any of the N items can be chosen to receive the most effort. Any of the remaining $N - 1$ can then be chosen to receive the second-most effort, and so on.

If $e_1 = \bar{e}$, then it is possible that she will also choose $e_2 = \bar{e}$, in which case the first inequality is weak. The equality of these efforts also reduces the number of possible optimal effort vectors. While there are still $N!$ permutations of I , several are associated with identical effort vectors and

therefore do not represent different combinations.

To understand Proposition 1, it is helpful to simplify the problem. Suppose that the actor is deciding between two choices, A and B , and suppose that the signal space is finite, with M possible signals: $\Psi = \{1, 2, \dots, M\}$. The actor's problem can be written

$$\max_{e_A, e_B} \sum_{i=1}^M \sum_{j=1}^M \Pr(s_A = i | e_A) \Pr(s_B = j | e_B) \max \{E(\theta_A | s_A = i, e_A), E(\theta_B | s_B = j, e_B)\} - c(e_A) - c(e_B).$$

With a little bit of oversimplification, we can provide clear intuition. Suppose that the actor chooses $e_A = e_B = e$. Whenever the signal for A is “higher” than the signal for B , she will choose A , and vice versa. What should the actor do when signals are identical? She can choose either A or B and receive the same payoff.

Suppose now that she increases effort on A and decreases effort on B a little bit, so that the total cost of effort is left unchanged. Her choice when the signals differ is still clear. In the cases in which the signal for A is higher, she may be a little better or worse off, just as in the cases in which the signal for B is higher. But the two sets of cases perfectly offset: the likelihood that the signal for A is i and the signal for B is j is identical to the likelihood that the signal for A is j and the signal for B is i . The increase (decrease) in the payoff in the first case is equal to the decrease (increase) in the payoff in the second case.

We are therefore interested in the cases in which the signals are identical. We can separate signals into two groups. The first group contains the (perhaps empty) set of signals for which $\frac{d}{de} E(\theta_x | s_x = i, e) = 0$. Within this set, a little more effort on A and a little less on B does not affect the choice between A and B when the signals are both i . The second group contains the (by assumption non-empty) set of signals for which $\frac{d}{de} E(\theta_x | s_x = i, e) \neq 0$. In this set, either $\frac{d}{de} E(\theta_x | s_x = i, e) > 0$ or $\frac{d}{de} E(\theta_x | s_x = i, e) < 0$. In either case, $\frac{d}{de} \max \{E(\theta_A | s_A = i, e + \varepsilon), E(\theta_B | s_B = i, e - \varepsilon)\}_{\varepsilon=0} > 0$. To see this, suppose that $E(\theta_x | s_x = i, e)$ is increasing in effort. Then item A will be chosen if signal i is observed for both A and B , and the expected payoff is higher than if $e_A = e_B = e$. If, instead, $E(\theta_x | s_x = i, e)$ is decreasing in effort, then item B will be chosen if signal i is observed for both A and B , and the expected payoff is again higher than if $e_A = e_B = e$.

The intuition above is hopefully relatively clear, but is not entirely accurate. We therefore

provide more precise intuition in the next five paragraphs.

Suppose that $e_A = e_B = e$ at the optimum and consider a very small change in efforts, increasing e_A to $e + \varepsilon$ and decreasing e_B to $e - \varepsilon$. This change will not have a first-order effect on cost because the marginal cost of the two efforts are equal. Will the expected benefit rise?

We can evaluate the derivative with respect to ε of each term in the objective. There are M^2 such terms. If we think of these terms as occupying an $M \times M$ matrix, then we can separate those terms to M on-diagonal and $M(M - 1)$ off-diagonal terms. Considering the off-diagonal terms first, let the two received signals be $i \neq j$. If the expected payout resulting from signals $s_A = i, s_B = j$ increases by $Z(\varepsilon)$, then the expected payout resulting from signals $s_A = j, s_B = i$ decreases by $Z(\varepsilon)$. Evaluated at $\varepsilon = 0$, the likelihood of each of those signal pairs is identical. Therefore, all off-diagonal terms cancel.

The important terms are the on-diagonal terms, that result when $s_A = s_B$. What is the actor to do when she observes identical signals? Each diagonal term of the objective can be written

$$\Pr(s_A = i | e + \varepsilon) \Pr(s_B = i | e - \varepsilon) \max \{E(\theta_A | s_A = i, e + \varepsilon), E(\theta_B | s_B = i, e - \varepsilon)\}.$$

To a first-order approximation, $\Pr(s_A = i | e + \varepsilon) \Pr(s_B = i | e - \varepsilon)$ does not vary with ε . More precisely,

$$\begin{aligned} & \frac{d}{d\varepsilon} [\Pr(s_A = i | e + \varepsilon) \Pr(s_B = i | e - \varepsilon)]_{\varepsilon=0} \\ &= \frac{d}{d\varepsilon} \Pr(s_A = i | e + \varepsilon) \Pr(s_B = i | e - \varepsilon)_{\varepsilon=0} \\ & \quad - \Pr(s_A = i | e + \varepsilon) \frac{d}{d\varepsilon} \Pr(s_B = i | e - \varepsilon)_{\varepsilon=0} \\ &= 0. \end{aligned}$$

Therefore, we are solely interested in $\frac{d}{d\varepsilon} \max \{E(\theta_A | s_A = i, e + \varepsilon), E(\theta_B | s_B = i, e - \varepsilon)\}_{\varepsilon=0}$. At $\varepsilon = 0$, $E(\theta_A | s_A = i, e + \varepsilon) = E(\theta_B | s_B = i, e - \varepsilon)$. If $\frac{d}{d\varepsilon} E(\theta_x | s_x = i, e + \varepsilon) = 0$ for all $\psi \in \{1, 2, \dots, M\}$, then the derivative of the objective is zero and $e_A = e_B = e$ could indeed be optimal. If, however, $\frac{d}{d\varepsilon} E(\theta_x | s_x = i, e + \varepsilon) \neq 0$ for at least one $\psi \in \{1, 2, \dots, M\}$ that occurs with positive probability, then the derivative of the objective is non-zero and $e_A = e_B = e$ is sub-optimal. By Assumption 3, equal efforts are therefore never optimal.

This completes our discussion of Proposition 1. It is optimal to allocate strictly different efforts to learning about each of the N items. We now turn to the consequences of this unequal decision rule.

3.2 Asymmetric selection likelihoods

Proposition 1 has a clear consequence in terms of the likelihood of a particular option being chosen. Recall that the actor's choice is denoted $y \in I$.

Remark 1. *Once an optimal effort vector is selected, it is generically the case that $\Pr(y = i) \neq \Pr(y = j)$, $i, j \in I$.*

Because the efforts all differ, each item has differing likelihoods of being associated with each signal and each signal's effect on the actor's posterior also differs by item. Therefore, the probability of a given item being chosen will not generally be the same across items.

This may be unclear, so we will briefly consider the two-signal example discussed above, and to which we return later. Suppose that there are N items and only two signals, $s_i \in \{L, H\}$, that satisfy

$$\Pr(s_i = H | \theta_i \geq \bar{\theta}) = \Pr(s_i = L | \theta_i < \bar{\theta}) = \frac{1 + e_i}{2},$$

and

$$\Pr(s_i = H | \theta_i < \bar{\theta}) = \Pr(s_i = L | \theta_i \geq \bar{\theta}) = \frac{1 - e_i}{2}.$$

Let $\Pr(\theta_i \geq \bar{\theta}) \equiv Q \in (0, 1)$. If $Q > 1/2$, then the likelihood of item 1 (the item toward which the most effort is directed) being chosen is simply the likelihood that item 1 receives a high signal. This is

$$\begin{aligned} \Pr(y = 1) &= Q \frac{1 + e_1}{2} + (1 - Q) \frac{1 - e_1}{2} \\ &= \frac{1}{2} + \left(Q - \frac{1}{2}\right) e_1 \\ &> \frac{1}{2}. \end{aligned}$$

So the likelihood of the first item being chosen is greater than $1/2$, even though there are N items! Clearly, the likelihood that any of the other $N - 1$ being chosen is much lower.

If $Q < 1/2$, it is not as easy to show that, typically, $\Pr(y = i) \neq \Pr(y = j)$, $i, j \in I$. If we narrow focus once again to a case with only two ($N = 2$) items in the option set, so that $I = \{A, B\}$, we show in Proposition 3 in Section 4 that there are two optimal effort vectors that solve the agent's problem: $\{e_A, e_B\} = \{(1-Q)QV, 0\}$ or $\{e_A, e_B\} = \{0, (1-Q)QV\}$.⁸ In each case, $\Pr(y = A) \neq \Pr(y = B)$ for $Q < 1/2$.

3.3 The effect of bias

Remark 1 suggests a way that even minor biases can cause major differences in outcomes for items in the choice set. The actor is indifferent to any of the $N!$ optimal effort vectors, but if she has even a small preference for one or more items, then she will naturally choose from the effort vectors that maximize the likelihood of that or those items being chosen.

For simplicity, in this section, we assume that the actor has a bias $b > 0$ in favor of item A . Her utility is

$$U = \theta_i - \sum_{j \in I} c(e_j) + 1_{y=A} \times b,$$

in which $1_{y=A}$ is an indicator taking a value of 1 if the actor chooses item A and zero otherwise. It follows almost immediately that she is no longer indifferent between the $N!$ effort vectors identified in Proposition 1. Once we introduce a bias, then she strictly prefers the subset of those vectors, which may be a singleton, that maximizes the likelihood of A being chosen.

For any positive bias, the optimal effort vector(s) differ(s) from any of the optimal vectors when the bias is zero, simply because the expected payoff for any vector is higher once the bias is introduced. As this is not particularly interesting – we could have introduced a bias of $-b$ if $y \neq A$ and the tension would be the same – we allow the bias to go to zero. In this case, the optimal effort vectors are a subset of the $N!$ vectors defined in Proposition 1.

Proposition 2. *As $b \rightarrow 0$, the actor's choice of e will maximize $\Pr(y = A)$ from the $N!$ optimal vectors identified in Proposition 1.*

Corollary 1. *[The Magnification of Small Biases] As $b \rightarrow 0$, the likelihood of item A being chosen is greater than $1/N$.*

⁸Note that, if $I = \{A, B\}$, then $N = 2$ so Proposition 1 has already established that there are $N! = 2$ optimal effort vectors.

Corollary 1 is the result that motivates the title of this manuscript: even infinitesimal biases lead to a difference in the likelihoods of the favored and disfavored items being chosen. As we established in Section 3.2, the difference can be large. If there are 100 items and the actor is infinitesimally biased in favor of one, her likelihood of choosing that one could be well over 1/2.

4 A simple example with two items, two types, and two signals

We began by showing that a very simple choice problem leads to asymmetric efforts toward learning about each of N possible choices. Those asymmetric efforts lead to asymmetric likelihoods of each item being selected. This simple decision theory problem and solution are interesting in their own rights: a simple symmetric choice problem yields an asymmetric decision. The problem becomes even more interesting once we allow for a small bias on the part of the decision-maker. The bias determines which of the optimal effort vectors the actor will select. If many actors have similar biases, then we have established that even very small biases can lead to large disparities in outcomes for members of different groups.

In this section, we focus on a simple example in which the number of items, the number of types, and the number of signals are all 2. This rather drastic simplification allows us to analytically solve for optimal efforts and the likelihoods of each item being chosen, and also allows us to delve more deeply into the relationship between the information environment and the magnification of small biases. The $2 \times 2 \times 2$ case also allows for considerably more clarity regarding the mechanisms discussed thus far.

4.1 The baseline with no bias

For now, *we return to the case with no bias*. There are two items indexed by $i \in \{A, B\}$, each with unknown type $\theta_i \in \{0, V\}$. The probability that $\theta_i = V$ is $Q \in (0, 1)$ for both items. As before, the agent exerts effort e_i to learn about item i , and this effort produces signals $s_i \in \{L, H\}$. These signals satisfy

$$\Pr(s_i = H | \theta_i = V) = \Pr(s_i = L | \theta_i = 0) = \frac{1 + e_i}{2},$$

and

$$\Pr(s_i = H | \theta_i = 0) = \Pr(s_i = L | \theta_i = V) = \frac{1 - e_i}{2}.$$

The effort cost is quadratic, so the actor's utility from choice y and efforts e_A, e_B is

$$U = \theta_y - \frac{1}{2} (e_A^2 + e_B^2).$$

So that the solution is interior (i.e., the optimal value of $e_i \in [0, 1)$), we make the following assumption:

Assumption 1: $V < \frac{1}{(1-Q)Q}$.

This $2 \times 2 \times 2$ version of the problem is simple enough that we can write the actor's posterior after observing either signal:

$$E(\theta_i | s_i = H, e_i) = \frac{\left(\frac{1+e_i}{2}\right) Q}{\left(\frac{1+e_i}{2}\right) Q + \left(\frac{1-e_i}{2}\right) (1-Q)} V, \quad (1)$$

$$E(\theta_i | s_i = L, e_i) = \frac{\left(\frac{1-e_i}{2}\right) Q}{\left(\frac{1-e_i}{2}\right) Q + \left(\frac{1+e_i}{2}\right) (1-Q)} V. \quad (2)$$

It follows immediately that:

Lemma 1. $E(\theta_i | s_i = H, e_i)$ is increasing in e_i and $E(\theta_i | s_i = L, e_i)$ is decreasing in e_i . That is, the more effort that the actor exerts toward learning, the more she updates in the "direction" of the signal.

There are four possible pairs of signals: (i) $s_A = H$ and $s_B = H$, (ii) $s_A = H$ and $s_B = L$, (iii) $s_A = L$ and $s_B = H$, and (iv) $s_A = L$ and $s_B = L$. The agent's problem is to:

$$\max_{e_A, e_B} \sum_{s_A \in \{L, H\}} \sum_{s_B \in \{L, H\}} \Pr(s_A | e_A) \Pr(s_B | e_B) \times \max\{E(\theta_A | s_A, e_A), E(\theta_B | s_B, e_B)\} - \frac{1}{2}(e_A^2 + e_B^2).$$

Because the priors for each item are identical, if one signal is high and one is low, the agent chooses the item that is associated with a high signal.

If the signals are both $s_A = s_B = L$ or $s_A = s_B = H$, then she chooses the item i that maximizes $E(\theta_i | s_i, e_i)$. It follows immediately from Lemma 1 that if the signals are both high, then she will

select the item toward which she exerted *greater* effort, and if the signals are both low, then she will select the item toward which she exerted *less* effort.

We now show that there are two effort pairs that maximize the agent's objective. The optimal effort pair is either $\{e_A, e_B\} = \{(1-Q)QV, 0\}$ or $\{e_A, e_B\} = \{0, (1-Q)QV\}$. Because the problem is *ex ante* symmetric, it is irrelevant whether more effort is put toward item *A* or item *B*. In either case, effort is put only toward one of the two.

Proposition 3. *In the $2 \times 2 \times 2$ special case, there are two optimal effort vectors that solve the agent's problem: $\{e_A, e_B\} = \{(1-Q)QV, 0\}$ or $\{e_A, e_B\} = \{0, (1-Q)QV\}$. In each case, the agent allocates strictly positive effort to one item and no effort to the other.*

Corollary 2. *If $\{e_A, e_B\} = \{(1-Q)QV, 0\}$, then the optimal choice is $y(s_A, s_B) = \begin{cases} A & \text{if } s_A = H \\ B & \text{if } s_A = L \end{cases}$*
and if $\{e_A, e_B\} = \{0, (1-Q)QV\}$, then the optimal choice is $y(s_A, s_B) = \begin{cases} A & \text{if } s_B = L \\ B & \text{if } s_B = H \end{cases}$.

Some intuition is in order. Suppose, arbitrarily, that $e_A > e_B$. If item *A* receives a high signal and item *B* a low signal, then *A* is chosen. If *A* and *B* both receive a high signal, then *A* will be chosen because effort e_A is higher, so the actor's posterior for *A* is higher. If item *B* receives a high signal and item *A* a low signal, then *B* is chosen. If both receive a low signal, then *B* is still chosen because effort e_B is lower, meaning item *B*'s posterior is higher (the actor updates downward less after seeing the signal because she invested less effort in it). Putting these four cases together, the actor chooses *A* if and only if *A* receives a high signal. The signal on item *B* is irrelevant to the actor's choice. Therefore, if $e_A > e_B$, the optimal effort toward item *B* is $e_B = 0$.

Similarly, if we arbitrarily assume that $e_B > e_A$, the same argument shows that e_A has positive marginal cost and no marginal benefit, so should be set to 0.

There is a third case in which $e_A = e_B$. However, by Proposition 1, this can never be optimal. Thus, there are two solutions to the actor's problem and both are characterized by an asymmetric effort choice.

As we have already shown, this argument extends to more than two options. When there are *N* *ex ante* identical items and one must be chosen, the optimal effort vector will feature *N* different effort choices, one of which may be zero. The assignment of efforts to items is arbitrary, but all

efforts will be different and one may be nil. It is critical to note that the fact that one of the items receives no effort is not particularly important for the results to follow. What is important is that *all of the efforts are different, as these differences lead to differential likelihoods of each item being chosen.*

Corollary 3. *If the prior is $Q > 1/2$, then the agent is more likely to choose the item toward which she invested more effort. If $Q < 1/2$, then the agent is more likely to choose the item toward which she invested less effort.*

This corollary is critical for understanding the results concerning bias that follow. If effort is put into learning about item i , then item i will be chosen if and only if the signal associated with that effort is positive. The likelihood of a good signal is greater than $1/2$ if and only if $Q > 1/2$, so item i will be more likely to be chosen if $Q > 1/2$ and less likely if $Q < 1/2$. Because of the simplicity of the $2 \times 2 \times 2$ setting, we can plot the underlying mathematics of Corollary 3 in Figure 3.

Suppose that $\{e_A, e_B\} = \{(1 - Q)QV, 0\}$ – called Solution 1 in Figure 3. Item A is chosen if and only if $s_A = H$. If $Q = 0$, then both items are always low quality and the agent exerts no effort toward learning about either. She can simply flip a coin in choosing an item. As Q increases, the effort she exerts toward item A increases. Because signals are more likely to be bad than good, her likelihood of choosing A decreases below $1/2$. As Q continues to increase, the likelihood that her signal is positive increases until, at $Q = 1/2$, her signal is equally likely to be positive and negative, so her likelihood of choosing A is $1/2$.

As Q increases above $1/2$, the likelihood of a high signal for item A continues to increase. Her effort toward learning about A , $e_A = (1 - Q)QV$, decreases, but the effect of a higher Q more than offsets the reduction in effort. As Q approaches 1, however, her effort goes to zero. She is indifferent between A and B , just as she was when $Q = 0$.

We can see that the two solutions are opposite. In one solution, effort is arbitrarily assigned to item A , in which case item A is more likely to be chosen when $Q > 1/2$ and item B is more likely to be chosen when $Q < 1/2$. In the other solution, effort is arbitrarily assigned to item B and the likelihoods of choosing item B follow the logic for item A in the first case. As in the more general model, the choice of effort vector is arbitrary.

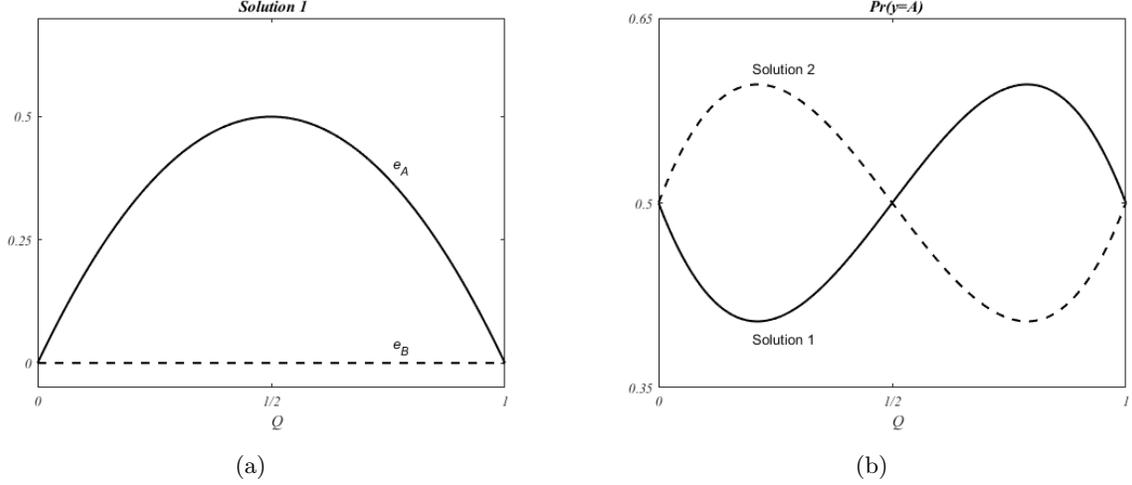


Figure 3: The two panels present the actor's effort choice (in Solution 1) and the probability of selecting item A (in both solutions) respectively. In the first panel, which corresponds to Solution 1, $e_A = (1 - Q)QV$ and $e_B = 0$. As Q increases towards $1/2$, the actor increases her total level of effort towards learning about item A because the uncertainty about ex ante quality is increasing. At $Q = 1/2$, uncertainty regarding ex ante quality is at its highest value and the actor's effort is maximized. Beyond $Q = 1/2$, the actor's effort decreases as it becomes more likely that a given item is high quality. In the second panel, the solid curve plots the likelihood that the actor chooses item A if $\{e_A, e_B\} = \{(1 - Q)QV, 0\}$. If $Q = 0$, then there is no reason to exert effort: both candidates are certainly bad. If $Q < 1/2$, the agent exerts more effort on candidate B so $\Pr(y = B) = \frac{1}{2} - \frac{e_B(1-2Q)}{2}$. As Q increases above 0, two things happen. First, mechanically, the likelihood of a good signal increases, increasing the likelihood of choosing B . Second, effort also increases, reducing the likelihood of choosing B . The second effect is more powerful for lower Q , and the first effect eventually overcomes the second as $Q \rightarrow 1/2$. At $Q = 1/2$, the likelihood of a good or bad signal equals $1/2$ regardless of effort, so the likelihood of the agent choosing candidate A also equals $1/2$ regardless of effort. At $Q = 1/2$, the signal is maximally informative, so effort is maximized. If $Q > 1/2$, the agent exerts more effort on candidate A so $\Pr(y = A) = \frac{1}{2} + \frac{e_B(2Q-1)}{2}$. As Q increases above $1/2$, two things happen. First, mechanically, the likelihood of a good signal increases, increasing the likelihood of choosing A . Second, effort also decreases, reducing the likelihood of choosing A . The first effect is more powerful for lower Q and the second effect eventually overcomes the first as $Q \rightarrow 1$. The dashed curve is simply the likelihood of choosing A under the alternative effort vector, $\{e_A, e_B\} = \{0, (1 - Q)QV\}$, and is therefore a reflection across the line $\Pr(y = A) = 1/2$.

4.2 The selection problem with bias

We now follow our outline from Section 3 by introducing a small bias. Because we are working within a much-simplified $2 \times 2 \times 2$ model, we can calculate optimal efforts and likelihoods of each item being chosen for non-zero biases. As before, we introduce a bias in favor of item A by adjusting the actor's utility function.

$$U = \theta_y - \frac{1}{2} (e_A^2 + e_B^2) + 1_{y=A} \times b.$$

We assume that the bias is sufficiently small for the problem to be interesting. That is, b is small enough that the actor's optimal efforts and choices results in her choosing item B if $s_B = H$ and $s_A = L$.

Proposition 4. *Given a small bias b :*

(i) *If $Q < 1/2$, then the optimal effort vector is $\{e_A, e_B\} = \{0, (1 - Q)QV + \frac{b(1-2Q)}{2}\}$ and the likelihood of choosing item A is,*

$$\Pr(y = A) = \frac{1}{2} + \frac{(1-2Q)(b(1-2Q)+2(1-Q)QV)}{4} > \frac{1}{2}.$$

(ii) *If $Q > 1/2$, then the optimal effort vector is $\{e_A, e_B\} = \{(1 - Q)QV + \frac{b(1-2Q)}{2}, 0\}$ and the likelihood of choosing item A is*

$$\Pr(y = A) = \frac{1}{2} + \frac{(2Q-1)(b(2Q-1)+2(1-Q)QV)}{4} > \frac{1}{2}.$$

(iii) *If $Q = 1/2$, then the optimal effort vector is either $\{e_A, e_B\} = \{0, (1 - Q)QV + \frac{b(1-2Q)}{2}\}$ or $\{e_A, e_B\} = \{(1 - Q)QV + \frac{b(1-2Q)}{2}, 0\}$ and the likelihood of choosing item A is $\Pr(y = A) = 1/2$.*

We can see the intuition for these results in Figure 3. The actor chooses the effort vector that maximizes the probability that she chooses item A. If $Q \in (0, 1/2)$, then she chooses “Solution 2”, in which effort is put toward item B. This is because the signal for item B is more likely to be bad. The probability that she chooses A is the dashed curve for $Q \in (0, 1/2)$. Similarly, if $Q \in (1/2, 1)$, then she chooses “Solution 1”, in which effort is put toward item A. This is because the signal for item A is more likely to be good. The probability that she chooses A is the solid curve for $Q \in (1/2, 1)$. Her likelihood of choosing A is the upper envelope of the curves in Figure 3, and exceeds $1/2$ for *almost all* Q .⁹

⁹If $b > 0$, it is not quite right that the optimal effort vector is either Solution 1 or Solution 2. Rather, it will be

We cannot, in the general model, establish how the magnification of bias interacts with the importance of the decision. There are several reasons, but one is that the type space for different items is allowed to be quite general. Now that we are focused on the $2 \times 2 \times 2$ model, we can see the interaction between the magnification and the importance of the decision simply by taking a comparative static with respect to V . The probability that item A is chosen is

$$\Pr(y = A) = \frac{1}{2} + |2Q - 1| \left(\frac{2(1 - Q)QV + |2Q - 1|b}{4} \right). \quad (3)$$

Note that, if $b \rightarrow 0$, this curve is the upper envelope in Figure 3. Regardless of whether $b \rightarrow 0$ or $b > 0$, this probability is increasing in V .

Corollary 4. [*The Glass Ceiling*] *If the prior $Q \neq 1/2$, then the likelihood of choosing item A is strictly increasing in the importance of the decision, V .*

The intuition is straightforward: increasing V increases the effort that the agent exerts to learn about item quality. Because efforts are always chosen so that A is the more likely item to be chosen, higher effort increases the likelihood that A is chosen.

Figure 4 plots the actor's effort choices and the likelihood of choosing item A for two values of V as a function of Q (with $b \rightarrow 0$). In the first panel, if $Q < 1/2$, the agent exerts effort to learning about item B and no effort to learning about item A . However, if $Q > 1/2$, the actor switches her choice such that she learns about item A but does not learn about item B . The solid curve represents the actor's effort choice if $V = 2$ and the dashed curve represents the actor's effort choice if $V = 2.5$. Notably, as V increases total effort also increases. In the second panel, the solid curve is the upper envelope of the curves in Figure 3 for $V = 2$ and the dashed curve plots the same for $V = 2.5$. As V increases, the probability of selecting item A increases.

For any value of Q other than the special cases of $Q = 0$, $Q = 1$, or $Q = 1/2$, higher values of V cause the agent to exert greater effort, which in turn causes a higher likelihood of choosing item A . This asymmetric effort, and therefore the asymmetric likelihood of choosing a particular item, would be present even without any bias. The bias does not directly cause this pattern. Instead, the bias selects the effort pair from two nearly equivalent options.

one of the two shown in Proposition 4. As the bias goes to zero, these solutions converge to Solution 1 and Solution 2 shown in Figure 3

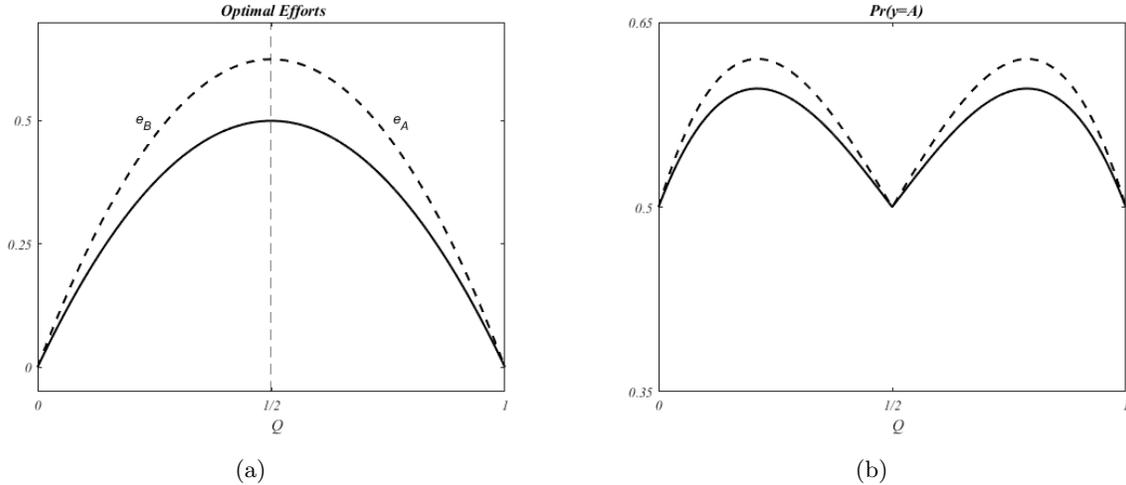


Figure 4: We plot the total efforts of the actor (the first panel) and the likelihood of choosing item A as a function of Q for $b \rightarrow 0$ (the second panel). For both values of V in the first panel, if $Q < 1/2$ the actor allocates effort to learning about item B but no effort to learning about item A . If $Q > 1/2$, the actor allocates effort to learning about item A but no effort to learning about item B . The solid curve represents the effort provided if $V = 2$ and the dashed curve represents the effort provided if $V = 2.5$. In the second panel, the solid curve is the upper envelope of the curves in Figure 3 for $V = 2$ and the dashed curve is the upper envelope for $V = 2.5$.

Figure 4 shows that even as the bias $b \rightarrow 0$, the likelihood of choosing the preferred item is strictly greater than $1/2$ except at three special points. If $Q = 0$ or $Q = 1$, choice is effectively irrelevant and the problem is degenerate. If $Q = 1/2$, the likelihood of a low or high signal is $1/2$ regardless of effort. This special case is a knife-edge result, and not generic.

Our *Glass Ceiling* corollary establishes that, in the simple $2 \times 2 \times 2$ model at least, more important decisions are met with greater effort but also a greater magnification of small biases. The former effect is precisely what we would intuit: if a decision is more important, then it deserves more attention. The latter effect is more surprising: if a decision is more important, then one would expect bias to be mitigated, not magnified!

We call Corollary 4 the Glass Ceiling corollary because it suggests that, as one looks at higher ranking positions in a corporate or political structure, we should expect fewer employees from disfavored groups. This theoretical prediction is consistent with many empirical studies of hierarchies.

4.3 Bias and the focus of investigations

Proposition 4 establishes the optimal effort pair in the $2 \times 2 \times 2$ model. Beyond discussing the fact that the pair is asymmetric and leads to asymmetric likelihoods of each item being chosen, we have not discussed the empirical prediction for each effort. There are two regions for Q that determine which effort vector the biased actor chooses.

Suppose that the actor is biased in favor of item A . If $Q > 1/2$, then signals are positive more often than they are negative for any effort greater than 0, so she directs effort toward learning about item A . She tends to find good news and therefore choose A . If $Q < 1/2$, then signals are negative more often, so she directs effort toward learning about item B . She will tend to find bad news and therefore choose A . Heads A wins, tails B loses!

In practice, what examples might fit these predictions? Suppose that there are two applicants for an apartments, both of whom are probably qualified renters. Landlords learn about the two candidates by performing due diligence. Most of the time, the background check will show no problems, but sometimes there would be a problem in which case the renter is disqualified. If the landlord is biased in favor of applicant A , then she will perform due diligence by calling his references. If no red flags arise, then applicant A will get the apartment. As red flags are not likely to arise, applicant A will be (maybe significantly) more likely to get the apartment.

A similar example concerns hiring. Suppose that there are two applicants for a job. If the job is difficult, then Q would be low. The hiring manager would investigate applicant B , against whom she is biased. More often than not, she will find negative information about B to disqualify her from the job, and A will be chosen.

The intuition here does not rely on the $2 \times 2 \times 2$ model or the fact that, in this model, one of the efforts is zero. In a more general model, as shown in Proposition 1, the actor exerts effort toward learning about each of the options, but exerts more toward some than others.

4.4 Implications for selected item quality

Our model has implications, as well, for the average item quality. These implications are important for distinguishing the model from alternative models of bias and choice, in particular as they relate to the glass ceiling. Because these examples relate to hiring, we will refer to “candidates” rather

than “items.”

Suppose that $Q > 1/2$, so $\{e_A, e_B\} = \left\{ \left(1 - Q\right)QV + \frac{b(1-2Q)}{2}, 0 \right\}$. Candidate B has expected quality of Q regardless of her signal, because the agent does not exert effort to learn about her ability. If candidate A is chosen, therefore, then her expected quality is greater than Q , and if she is not chosen, then her expected quality is less than Q .

Similarly, suppose that $Q < 1/2$, so $\{e_A, e_B\} = \left\{ 0, \left(1 - Q\right)QV + \frac{b(1-2Q)}{2} \right\}$. Candidate A has expected quality Q regardless of her signal. If candidate B is chosen, therefore, then her expected quality is greater than Q , and if she is not chosen, then her expected quality is less than Q .

These facts imply the following result.

Proposition 5. *If the actor exerts effort toward learning about candidate i , then candidate i is of higher expected quality if she is chosen than is candidate j if she is chosen. Candidate i is of lower expected quality if she is not chosen than is candidate j if she is not chosen. Formally:*

- (i) *If $Q > 1/2$, then $\Pr(\theta_A = V \mid y = A) > \Pr(\theta_B = V \mid y = B) = \Pr(\theta_B = V \mid y = A) > \Pr(\theta_A = V \mid y = B)$.*
- (ii) *If $Q < 1/2$, then $\Pr(\theta_B = V \mid y = B) > \Pr(\theta_A = V \mid y = A) = \Pr(\theta_A = V \mid y = B) > \Pr(\theta_B = V \mid y = A)$.*

These predictions differ from what one would get in a simple alternative model of biased choice. Consider a “threshold” model in which talent is normally distributed with identical means and variances for candidates A and B . The actor has utility equal to the hire’s talent, plus some additional benefit if she hires candidate A . In this model, the agent essentially has a lower bar for hiring the favored candidate. The expected quality of a new hire is therefore lower if the hire is part of the favored group. In addition, the expected quality of a rejected applicant is also lower if the candidate is from the favored group.

Our model and the threshold model have clear empirical implications. For example, Hebert (2019) studies the success of entrepreneurs who receive equity-based venture funding in France. She finds that entrepreneurs from the disfavored group perform better after receiving funding than entrepreneurs from the favored group. If we reasonably assume that fewer than half of entrepreneurs seeking funding are of high quality, then this result is consistent with our model. Future research

could distinguish between the validity of the threshold model and our model by comparing the performance of entrepreneurs who fail to receive funding. The threshold model predicts that members of the disfavored group will outperform conditional on non-funding. Our model predicts the opposite.

Our model also has implications for the ex post justification of a choice. If $Q > 1/2$, the actor allocates her effort to learning about the preferred candidate because she is more likely to uncover positive information. This implies that, on average, selected candidates will perform well when $Q > 1/2$. Therefore, if an actor dismisses accusations of bias based on the *realized* performance of the selected candidate (i.e., “while candidate A is from the favored group, candidate A ’s performance is superior. Candidate A was the best choice.”), such an argument is not necessarily valid. That is, our model shows that ex post success by a selected candidate does not imply the absence of bias.

5 Conclusion

We have analyzed the optimal behavior for an actor who must choose a single item from N *ex ante* identical options. She can exert effort to learn about the quality of each item before making her choice. She may have a bias in favor of one the items, but the bias is small relative to the importance of choosing the best item. This very simple problem results in three surprising findings.

First, we show that the actor’s optimal effort choice is asymmetric: if she is unbiased, then she exerts substantial effort on one arbitrarily chosen item, less on another, still less on a third, etc. This means that, once she has decided upon the assignment of efforts to items, her likelihood of choosing each item is not $1/N$. This asymmetric effort solution to a simple, symmetric decision problem is surprising.

Second, if the actor has a small bias in favor of one of the options, this removes the arbitrariness of the assignment of efforts to items: she will choose an assignment that maximizes the likelihood of her favored item being chosen. When learning is more likely to result in negative information, she will invest in learning about items against which she is biased. When learning is more likely to result in positive information, she will invest in learning about items toward which she is biased. In both cases, a favored item is more likely to be chosen, even if the bias is small. This means that

the likelihood of a favored item being chosen is above $1/N$, and sometimes well above 50% even if N is large. This effect is a *magnification of an infinitesimal bias*.

Third, we present a very simple example that allows us to solve, rather than simply characterize the solution to, the model. In this example, we present an especially surprising result. As the importance of making the right choice increases, the magnification of small biases increases as well. This is important, because one might expect that bias magnification would be mitigated by high stakes and by the opportunity to learn about quality. We show that, at least in our model, learning and stakes do not necessarily mitigate the problem. In fact, stakes can exacerbate it. If we consider the application of hiring, the model predicts that, as one moves up the corporate ladder into increasingly important roles within the firm, hiring disparities will increase. *A glass ceiling will appear*.

Our simple choice problem arises in many settings. For example, an activist investor may have increasing returns to investment in one particular stock if it is easier to force change with 20% of one firm's shares than 10% of shares in two different firms. In this case, the activist must choose one of N . A consumer must choose one car insurer. A student must choose one college. A firm must hire one CEO. In each case, the choice is one of N . We predict that in all of these cases, the actor will put substantial effort into learning about one option, less into the next, etc.

In many examples, especially concerning hiring, bias is arguably important for the outcomes we see in practice. Bertrand and Mullinathan (2004) find that applicants with black sounding names are less likely to receive replies when they submit resumes for job openings. What do we make of this within the context of the model? First, it suggests a bias against black applicants. Second, it does not imply anything about the magnitude of the bias – even small biases can lead to large differences in call-back and hiring rates. Third, it suggests that Q is high, so hiring managers are targeting effort toward learning about favored applicants. Assuming that most of the openings in the study were for lower level jobs, which is likely, given the way that workers were solicited, this appears consistent with the setting.

Other examples of racial and gender bias in hiring and equity investment abound. We have already discussed in the introduction the evidence in Hebert (2019) concerning equity investments in female- and male-led firms. Hebert (2019) shows that women (men) are more likely to receive equity investment in female (male) dominated fields. Assuming a context-specific bias – i.e., that

venture capitalists are biased in favor of the sex that more commonly starts firms in an industry – this is consistent with our model. She also finds that the performance of female (male) led start-ups is higher in male (female) dominated industries. Our model is consistent with these empirical findings, as well.

We do not model an extended setting in which the actor must choose $n < N$ options rather than just one. We believe that the intuition of our results would hold. The reason for asymmetric efforts is that efforts are compliments or substitutes when one is limited in how many items one can choose. More effort toward learning about one item can, for example, increase the likelihood that item is chosen, lessening the expected return to learning about another item. It can also decrease the likelihood that item is chosen, increasing the expected return to learning about another item. When one is unrestricted and can choose any number (including none) of the N items, this endogenous complementarity or substitutability does not arise. It is reasonably clear that, if the actor must choose $n < N$ of the items, the intuition behind our results is present.

If we allow that the intuition applies to the choice of $n < N$ items, then implications for finance are even more common. For example, individual investors will tend to own a limited number of individual stocks. Assuming that an investor wishes to own n different stocks, her problem is to choose the best n after researching her options. We predict that she will not research all possible companies equally, instead investing substantial time learning about some and less learning about others. Our model can therefore rationalize home bias as well as many other biases in investors' portfolios. For example, our model can explain why many investors' portfolios are heavily tilted towards socially responsible (SRI) stocks and *green* investment funds. If socially-minded agents are even infinitesimally biased in favor of SRI investments, then their research efforts will result in them choosing socially-conscious stocks far more often than other investment options. When many items or assets are *ex ante* identical in terms of payoffs – as should be the case in an efficient market – small biases are greatly magnified.

6 References

Bartos, Vojtech, Michal Bauer, Julie Chytilova, and Filip Matejka (2016). Attention discrimination: theory and field experiments with monitoring information acquisition. *American Economic Review*

106 (6), pp. 1437-1475.

Becker, Gary (1971). *The economics of discrimination*. 2nd Edition. *Chicago: University of Chicago Press*.

Bergemann, Dirk, and Juuso Valimäki (2006). *Bandit Problems*. *Helsinki Center of Economic Research Discussion Paper No. 93*.

Bertrand, Marianne, and Sendhil Mullinathan (2004). Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination. *American Economic Review*, 94 (4), pp. 991-1013.

Bertrand (2011). New perspectives on gender. *Handbook of Labor Economics*, 4, pp. 1543-1590

Blau, Francine, and Lawrence Kahn (2017). The gender wage gap: Extent, trends, and explanations. *Journal of Economic Literature* 55 (3), pp. 789-865.

Bordalo, Pedro, Katherine Coffman, Nicola Gennaioli, and Andrei Shleifer (2016). Stereotypes. *The Quarterly Journal of Economics*, 131 (3), pp. 1753-1794.

Gennaioli, Nicola, and Andrei Shleifer (2010). What comes to mind? *The Quarterly Journal of Economics*, 125 (4), pp. 1399-1433.

Goldin, Claudia, and Cecilia Rouse (2000). Orchestrating impartiality: The impact of “blind” auditions on female musicians. *American Economic Review*, 90 (4), pp. 715-741.

Hebert, Camille (2018). *Mind the gap: Gender stereotypes and entrepreneur financing*. Working paper.

Kahneman, Daniel (2011). *Thinking Fast and Slow*. 1st Edition. New York: *Farrar, Straus, and Giroux*.

Rothschild, Michael (1974). A Two-Armed Bandit Theory of Market Pricing. *Journal of Economic Theory*, 9, pp. 185-202.

van Nieuwerburgh, Stijn, and Laura Veldkamp (2009). Information immobility and the home bias puzzle. *Journal of Finance*, 64 (3), pp. 1187-1215.

van Nieuwerburgh, Stijn, and Laura Veldkamp (2010). Information acquisition and under-diversification. *The Review of Economic Studies* 77 (2), pp. 779-805

7 Appendix A: Proofs omitted from main text

Proof of Proposition 1. For clarity, we provide a proof when $N = 2$ and the set of signals $\Psi \equiv \{1, 2, \dots, M\}$. We discuss at the end how the proof can be extended to $N > 2$. Since $N = 2$, there are two items given by $I \in \{A, B\}$. Define $g_m(e) = \Pr(s_i = m | e_i = e)$ for $i \in I$ and $m \in \Psi$. The actor's problem is

$$U^* = \max_{e_A e_B} U(e_A, e_B) = \max_{e_A e_B} \sum_{m'=1}^M \sum_{m=1}^M g_{m'}(e_A) g_m(e_B) \max\{E(\theta|m', e_A), E(\theta|m, e_B)\} - c(e_A) - c(e_B). \quad (4)$$

Now suppose for purposes of a contradiction that the agent's optimal choice is $e_A^* = e_B^* = e$. Consider the following deviation $e'_A = e + \varepsilon$ and $e'_B = e - \varepsilon$. The actor's payoff under this deviation is

$$U(e + \varepsilon, e - \varepsilon) = \sum_{m'=1}^M \sum_{m=1}^M g_{m'}(e + \varepsilon) g_m(e - \varepsilon) \max\{E(\theta|m', e + \varepsilon), E(\theta|m, e - \varepsilon)\} - c(e + \varepsilon) - c(e - \varepsilon)$$

Differentiating this expression with respect to ε gives

$$\begin{aligned} \frac{d}{d\varepsilon} U(e + \varepsilon, e - \varepsilon) = & \sum_{m'=1}^M \sum_{m=1}^M (-g_{m'}(e + \varepsilon) g'_m(e - \varepsilon) + g'_{m'}(e + \varepsilon) g_m(e - \varepsilon)) \max\{E(\theta|m', e + \varepsilon), E(\theta|m, e - \varepsilon)\} \\ & + \sum_{m'=1}^M \sum_{m=1}^M g_{m'}(e + \varepsilon) g_m(e - \varepsilon) \frac{d}{d\varepsilon} \max\{E(\theta|m', e + \varepsilon), E(\theta|m, e - \varepsilon)\} \\ & - c'(e + \varepsilon) + c'(e - \varepsilon) \quad (5) \end{aligned}$$

And evaluating this derivative at $\varepsilon = 0$ gives

$$\frac{d}{d\varepsilon}U(e + \varepsilon, e - \varepsilon)|_{\varepsilon=0} = \sum_{m=1}^M (g_m(e))^2 \frac{d}{d\varepsilon} \max\{E(\theta|m, e + \varepsilon), E(\theta|m, e - \varepsilon)\}|_{\varepsilon=0} > 0 \quad (6)$$

Some explanation is in order before evaluating further. To see the intuition for this result, consider two signals $\{j, k\} \in \{1, \dots, M\}$, and consider the case when $s_A \in \{j, k\}$ and $s_B \in \{j, k\}$. We will now explore the three lines from equation (5) with respect to these (and only these) pairs of signals. The first line of equation (5) is written as

$$\begin{aligned} & (-g_j(e + \varepsilon)g'_j(e - \varepsilon) + g'_j(e + \varepsilon)g_j(e - \varepsilon)) \max\{E(\theta|j, e + \varepsilon), E(\theta|j, e - \varepsilon)\} \\ & + (-g_j(e + \varepsilon)g'_k(e - \varepsilon) + g'_j(e + \varepsilon)g_k(e - \varepsilon)) \max\{E(\theta|j, e + \varepsilon), E(\theta|k, e - \varepsilon)\} \\ & + (-g_k(e + \varepsilon)g'_j(e - \varepsilon) + g'_k(e + \varepsilon)g_j(e - \varepsilon)) \max\{E(\theta|k, e + \varepsilon), E(\theta|j, e - \varepsilon)\} \\ & + (-g_k(e + \varepsilon)g'_k(e - \varepsilon) + g'_k(e + \varepsilon)g_k(e - \varepsilon)) \max\{E(\theta|k, e + \varepsilon), E(\theta|k, e - \varepsilon)\} \end{aligned}$$

Now, evaluating these terms at $\varepsilon = 0$ gives

$$\begin{aligned} & (-g_j(e)g'_j(e) + g'_j(e)g_j(e)) \max\{E(\theta|j, e), E(\theta|j, e)\} \\ & + (-g_j(e)g'_k(e) + g'_j(e)g_k(e)) \max\{E(\theta|j, e), E(\theta|k, e)\} \\ & + (-g_k(e)g'_j(e) + g'_k(e)g_j(e)) \max\{E(\theta|k, e), E(\theta|j, e)\} \\ & + (-g_k(e)g'_k(e) + g'_k(e)g_k(e)) \max\{E(\theta|k, e), E(\theta|k, e)\} = 0 \end{aligned}$$

It is straightforward to see that the first and fourth lines are equal to zero. In addition, the second and third lines sum together to equal zero. We now turn to the second set of terms in equation (5). Suppose without loss of generality that $E[\theta|j, e] > E[\theta|k, e]$. Then for ε small enough and again

restricting attention to only $s_A \in \{j, k\}$ and $s_B \in \{j, k\}$, these terms are written as

$$\begin{aligned}
& g_j(e + \varepsilon)g_j(e - \varepsilon) \frac{d}{d\varepsilon} \max\{E(\theta|j, e + \varepsilon), E(\theta|j, e - \varepsilon)\} \\
& + g_j(e + \varepsilon)g_k(e - \varepsilon) \frac{d}{de} E(\theta|j, e + \varepsilon) \\
& + g_k(e + \varepsilon)g_j(e - \varepsilon) \left(-\frac{d}{de} E(\theta|j, e - \varepsilon) \right) \\
& + g_k(e + \varepsilon)g_k(e - \varepsilon) \frac{d}{d\varepsilon} \max\{E(\theta|k, e + \varepsilon), E(\theta|k, e - \varepsilon)\}
\end{aligned}$$

And, evaluating these terms at $\varepsilon = 0$ gives

$$\begin{aligned}
& g_j(e)g_j(e) \frac{d}{d\varepsilon} \max\{E(\theta|j, e + \varepsilon), E(\theta|j, e - \varepsilon)\}|_{\varepsilon=0} \\
& + g_j(e)g_k(e) \frac{d}{de} E(\theta|j, e) \\
& + g_k(e)g_j(e) \left(-\frac{d}{de} E(\theta|j, e) \right) \\
& + g_k(e)g_k(e) \frac{d}{d\varepsilon} \max\{E(\theta|k, e + \varepsilon), E(\theta|k, e - \varepsilon)\}|_{\varepsilon=0}
\end{aligned}$$

Again, it is straight forward to see that the second and third lines sum to zero, so that the only terms which remain satisfy $s_A = s_B$. And finally, it is clear that $-c'(e + \varepsilon) + c'(e - \varepsilon)|_{\varepsilon=0} = 0$. The preceding arguments establish that

$$\frac{d}{d\varepsilon} U(e + \varepsilon, e - \varepsilon)|_{\varepsilon=0} = \sum_{m=1}^M (g_m(e))^2 \frac{d}{d\varepsilon} \max\{E(\theta|m, e + \varepsilon), E(\theta|m, e - \varepsilon)\}|_{\varepsilon=0}$$

We now need to show that $\frac{d}{d\varepsilon} U(e + \varepsilon, e - \varepsilon)|_{\varepsilon=0} > 0$. When the signals are equal (that is, $s_A = s_B$), the derivative of the maximum may not exist at $\varepsilon = 0$. For our purposes now, we will consider instead the right derivative, since this coincides with the limit as $\varepsilon > 0$ goes to 0. Note that the right derivative with respect to ε corresponds to the right derivative of e'_A and the left derivative of e'_B , since e'_A is linearly increasing in ε while e'_B is linearly decreasing in ε . Now, by assumption, there exists some signal $m \in \{1, \dots, M\}$ such that $\frac{d}{de} E(\theta|m, e) \neq 0$ and $g_m(e) > 0$. Suppose first that $\frac{d}{de} E(\theta|m, e) > 0$. Then,

$$(g_m(e))^2 \frac{d^+}{d\varepsilon} \max\{E(\theta|m, e + \varepsilon), E(\theta|m, e - \varepsilon)\}|_{\varepsilon=0} = (g_m(e))^2 \frac{d}{de} E(\theta|m, e) > 0$$

since $E(\theta|m, e + \varepsilon) > E(\theta|m, e - \varepsilon)$ for ε small enough. Now suppose that $\frac{d}{de}E(\theta|m, e) < 0$. Then,

$$(g_m(e))^2 \frac{d^+}{d\varepsilon} \max\{E(\theta|m, e + \varepsilon), E(\theta|m, e - \varepsilon)\}|_{\varepsilon=0} = (g_m(e))^2 \left(-\frac{d}{de}E(\theta|m, e) \right) > 0$$

since $E(\theta|m, e + \varepsilon) < E(\theta|m, e - \varepsilon)$ for ε small enough. Finally, suppose that $\frac{d}{de}E(\theta|m, e) = 0$.

Then

$$(g_m(e))^2 \frac{d^+}{d\varepsilon} \max\{E(\theta|m, e + \varepsilon), E(\theta|m, e - \varepsilon)\}|_{\varepsilon=0} = 0$$

Thus every term in equation (6) is weakly positive with at least one term strictly positive by assumption, and this implies that $\frac{d}{d\varepsilon}U(e + \varepsilon, e - \varepsilon)|_{\varepsilon=0} > 0$. Since the actor's problem is continuous differentiable with respect to e_A and e_B , this implies that there exists some $\bar{\varepsilon}$ small enough such that $U(e + \bar{\varepsilon}, e - \bar{\varepsilon}) > U(e, e)$ which contradicts the optimality of $e_A^* = e_B^* = e$. Extending the proof to $N > 2$ is straight forward. The same argument implies that if any two efforts are equal $e_i^* = e_l^* = e$, then a deviation of the form $e'_i = e + \varepsilon$ and $e'_l = e - \varepsilon$ necessarily increases the actor's payoff for ε small enough. The only complication relative to the case with $N = 2$ is that the signal m for which $\frac{d}{de}E(\theta|m, e) \neq 0$ must be payoff relevant. That is, given all other efforts e_n for $n \in I/i, l$ and all possible signals $j \in \{1, \dots, M\}$ there must exist a signal m such that $E(\theta|m, e) = \max\{E(\theta|m, e), E(\theta|j, e_n)\}$. Extending the proof to a non-finite signal space is also straight forward. Recall though that the signal space cannot be purely continuous, since by assumption there exists some signal m such that $g_m(e) > 0$. Any signals which occur with zero probability are second order for small changes in ε . The proof therefore extends as is to only those signals which occur with positive probability. \square

Proof of Proposition 2. Define $e_0^* \subset X^N$ to be the set of optimal efforts when $b = 0$, and define $e_+^* \subset X^N$ to be the set of optimal effort vectors in the limit as $b \rightarrow 0$. Since the agent's payoff is continuous with respect to b , the set of efforts $e^* \subset X^N$ is upper hemi continuous by the maximum theorem and as a result $e_+^* \subseteq e_0^*$. Now for purposes of a contradiction, suppose that some element $e_{+,j}^* \in e_+^* \subseteq e_0^*$ does not maximize $\Pr(y = A)$. That is, suppose that there exists some solution $e_{0,k}^* \in e_0^*$ such that $\Pr(y = A|e = e_{0,k}^*) > \Pr(y = A|e = e_{+,j}^*)$. Denote the agent's maximized payoff

by

$$U^* = \max_{e \in X^N} U(e) = \max_{e_A e_B} \sum_{m=1}^M \dots \sum_{m=1}^M \prod_{i \in I} g_m(e_i) \max\{E(\theta|m, e_A), E(\theta|m, e_B), \dots\} \\ - \sum_{i \in I} c(e_i) + \Pr(y = A|e)b.$$

Then,

$$\frac{dU^*}{db} = \frac{dU}{db} + \frac{dU}{de} \frac{de}{db} = \Pr(y = A|e^*),$$

since at an optimum $\frac{dU}{de} = 0$ by concavity. So, for a sufficiently small increase in b , the agent's maximized payoff increases by the probability that the agent selects item A . As a result, if $\Pr(y = A|e = e_{0,k}^*) > \Pr(y = A|e = e_{+,j}^*)$, then in the limit as $b \rightarrow 0$, $U(e_{0,k}^*) > U(e_{+,j}^*)$ which contradicts the optimality of $e_{+,j}^*$. \square

Proof of Corollary 1. A naive decision rule of picking one item at random leads to a uniform probability of $1/N$ for any item being chosen, including item A . The naive decision rule requires no allocation of effort. Thus, strategic allocation of costly effort across choices to maximize $\Pr(y = A|e)$ implies that the likelihood of choosing item A is greater than $1/N$. \square

Proof of Lemma 1. Simplifying equations (1) and (2) gives,

$$E(\theta_i | s_i = H, e_i) = \frac{(1 + e_i)QV}{1 + (2Q - 1)e_i} \quad (7)$$

$$E(\theta_i | s_i = L, e_i) = \frac{(1 - e_i)QV}{1 - (2Q - 1)e_i}, \quad (8)$$

and differentiating with respect to e_i gives,

$$\frac{d}{de_i} E(\theta_i | s_i = H, e_i) = \frac{2(1 - Q)QV}{(1 + (2Q - 1)e_i)^2} > 0 \quad (9)$$

$$\frac{d}{de_i} E(\theta_i | s_i = L, e_i) = \frac{-2(1 - Q)QV}{(1 - (2Q - 1)e_i)^2} < 0. \quad (10)$$

\square

Proof of Proposition 3. By Proposition 1, any solution features $e_i > e_j$. First, consider the case in which $e_A > e_B = 0$. The agent chooses item A when $s_A = s_B = H$ and chooses item B when

$s_A = s_B = L$. The agent's problem is given by,

$$\begin{aligned} \max_{e_A, e_B} & \Pr(s_A = H|e_A)E(\theta_A|s_A = H, e_A) + \Pr(s_A = L|e_A) \left(\Pr(s_B = H|e_B)E(\theta_B|s_B = H, e_B) \right. \\ & \left. + \Pr(s_B = L|e_B)E(\theta_B|s_B = L, e_B) \right) - \frac{e_A^2}{2} - \frac{e_B^2}{2}, \end{aligned} \quad (11)$$

which may be rewritten as,

$$\max_{e_A, e_B} \frac{1 + e_A}{2} QV + \left((1 - Q) \frac{1 + e_A}{2} + Q \frac{1 - e_A}{2} \right) QV - \frac{e_A^2}{2} - \frac{e_B^2}{2}.$$

The objective function is concave in e_A and e_B and first-order conditions are necessary and sufficient to solve for an internal solution. First-order conditions with respect to e_A and e_B yield,

$$0 = (1 - Q)QV - e_A \quad (12)$$

$$0 = -e_B, \quad (13)$$

yielding the solution $\{e_A, e_B\} = \{(1 - Q)QV, 0\}$. The case in which $e_B > e_A = 0$ is solved similarly and yields the solution $\{e_A, e_B\} = \{0, (1 - Q)QV\}$. Under both solutions, the agent's expected utility is given by,

$$E(U) = QV + \frac{(1 - Q)^2 Q^2 V^2}{2}. \quad (14)$$

□

Proof of Corollary 2. Note that $E(\theta_i|s = H, e_i = 0) = E(\theta_i|s = L, e_i = 0) = E(\theta_i) = QV$. When the actor exerts no effort learning about item i , her signal is completely uninformative, and her posterior is unchanged regardless of the realization of the signal. Now consider the case when the actor chooses the solution $\{e_A, e_B\} = \{(1 - Q)QV, 0\}$. Since $\frac{d}{de_i} E(\theta_i|s_i = H, e_i) > 0$ and $\frac{d}{de_i} E(\theta_i|s_i = L, e_i) < 0$, we have

$$E(\theta_A|s_A = H, e_A = (1 - Q)QV) > E(\theta_B|s_B = H, e_B = 0) = QV, \quad (15)$$

$$QV = E(\theta_B|s_B = L, e_B = 0) > E(\theta_A|s_A = L, e_A = (1 - Q)QV). \quad (16)$$

Since $E(\theta_A|s_A = H, e_A = (1 - Q)QV) > QV$, the actor chooses $y = A$ whenever $s_A = H$, and since

$E(\theta_A|s_A = L, e_A = (1 - Q)QV) < QV$, the actor chooses $y = B$ whenever $s_A = L$. The alternative case with $\{e_A, e_B\} = \{0, (1 - Q)QV\}$ follows from the same argument. \square

Proof of Proposition 4. There are multiple cases to consider. First, consider the case in which $s_A = s_B = H$. The agent's optimal choice is candidate B if, and only if, the condition,

$$E(\theta_B|s_B = H, e_B) - E(\theta_A|s_A = H, e_A) > b, \quad (17)$$

holds. If the condition in (17) does hold, then the agent's optimal choice is candidate B if $s_A = L$ and $s_B = H$, is candidate A if $s_A = H$ and $s_B = L$, and is candidate A if $s_A = s_B = L$. Thus, if $E(\theta_B|s_B = H, e_B) - E(\theta_A|s_A = H, e_A) > b$, the agent's problem is given by,

$$\begin{aligned} \max_{e_A, e_B} & \Pr(s_B = H|e_B)E(\theta_B|s_B = H, e_B) + \Pr(s_B = L|e_B) \left(\Pr(s_A = H|e_A)E(\theta_A|s_A = H, e_A) \right. \\ & \left. + \Pr(s_A = L|e_A)E(\theta_A|s_A = L, e_A) + b \right) - \frac{e_A^2}{2} - \frac{e_B^2}{2}, \end{aligned} \quad (18)$$

which may be rewritten as,

$$\max_{e_A, e_B} \frac{1 + e_B}{2} QV + \left((1 - Q) \frac{1 + e_B}{2} + Q \frac{1 - e_B}{2} \right) (QV + b) - \frac{e_A^2}{2} - \frac{e_B^2}{2}.$$

The objective function is concave in e_A and e_B and first-order conditions are necessary and sufficient to solve for an internal solution. First-order conditions with respect to e_A and e_B yield,

$$0 = -e_A \quad (19)$$

$$0 = (1 - Q)QV + \frac{b(1 - 2Q)}{2} - e_B, \quad (20)$$

yielding the solution $\{e_A, e_B\} = \left\{ 0, (1 - Q)QV + \frac{b(1 - 2Q)}{2} \right\}$. It may be quickly verified that the condition in (17) is satisfied for a sufficiently small b ,

$$E(\theta_B|s_B = H, e_B) - E(\theta_A|s_A = H, e_A) - b = \left(\frac{Q \frac{1 + e_B}{2}}{Q \frac{1 + e_B}{2} + (1 - Q) \frac{1 - e_B}{2}} - Q \right) V - b \quad (21)$$

$$= \frac{Q(1 - Q)V}{\frac{1}{2e_B} + \left(Q - \frac{1}{2}\right)} - b, \quad (22)$$

which is strictly positive if, and only if,¹⁰

$$b < b^* \equiv \frac{1 - 2(1 - Q)Q(1 - 2Q)V + \sqrt{1 + 4(1 - Q)Q(2Q - 1)V}}{(2Q - 1)^2}. \quad (23)$$

The likelihood of choosing candidate A with $\{e_A, e_B\} = \left\{0, (1 - Q)QV + \frac{b(1 - 2Q)}{2}\right\}$ is,

$$\Pr(s_B = L|e_B) = Q \frac{1 - e_B}{2} + (1 - Q) \frac{1 + e_B}{2} \quad (24)$$

$$= Q \frac{1 - (1 - Q)QV - \frac{b(1 - 2Q)}{2}}{2} + (1 - Q) \frac{1 + (1 - Q)QV + \frac{b(1 - 2Q)}{2}}{2} \quad (25)$$

$$= \frac{1}{2} + \frac{(1 - 2Q)(b(1 - 2Q) + 2(1 - Q)QV)}{4}, \quad (26)$$

which is strictly smaller than $\frac{1}{2}$ if $Q > \frac{1}{2}$, equals $\frac{1}{2}$ at $Q = \frac{1}{2}$, and is greater than $\frac{1}{2}$ if $Q < \frac{1}{2}$. The likelihood of choosing candidate B is,

$$\Pr(s_B = H|e_B) = \frac{1}{2} - \frac{(1 - 2Q)(b(1 - 2Q) + 2(1 - Q)QV)}{4}.$$

Furthermore, the agent's expected utility is given by,

$$E(U) = QV + \frac{b}{2} + \frac{\left((1 - Q)QV + \frac{(1 - 2Q)b}{2}\right)^2}{2}, \quad (27)$$

and the limit of $E(U)$ as $b \rightarrow 0$ is given by,

$$E(U) = QV + \frac{(1 - Q)^2 Q^2 V^2}{2}, \quad (28)$$

Next, consider when the condition in (17) does not hold and consider the case in which $s_A = s_B = L$.

The agent's optimal choice is candidate B if, and only if, the condition,

$$E(\theta_B|s_B = L, e_B) - E(\theta_A|s_A = L, e_A) > b, \quad (29)$$

holds. If the condition in (17) does not hold and the condition in (29) does hold, then the agent's

¹⁰Technically, there are two solutions for b^* . However, one solution of b^* can be ruled out if $e_B > 0$ (which only occurs if $Q \leq \frac{1}{2}$ as is shown shortly), because it is weakly negative valued.

optimal choice is candidate A if $s_A = s_B = H$, is candidate A if $s_A = H$ and $s_B = L$, and is candidate B if $s_A = L$ and $s_B = H$. Thus, if $E(\theta_B|s_B = H, e_B) - E(\theta_A|s_A = H, e_A) \leq b$, and $E(\theta_B|s_B = L, e_B) - E(\theta_A|s_A = L, e_A) > b$, the agent's problem is given by,

$$\begin{aligned} \max_{e_A, e_B} & \Pr(s_A = H|e_A) (E(\theta_A|s_A = H, e_A) + b) \\ & + \Pr(s_A = L|e_A) \left(\Pr(s_B = H|e_B) E(\theta_B|s_B = H, e_B) \right. \\ & \left. + \Pr(s_B = L|e_B) E(\theta_B|s_B = L, e_B) \right) - \frac{e_A^2}{2} - \frac{e_B^2}{2}, \end{aligned} \quad (30)$$

which may be rewritten as,

$$\begin{aligned} \max_{e_A, e_B} & \frac{1+e_A}{2} QV + \left(Q \frac{1+e_A}{2} + (1-Q) \frac{1-e_A}{2} \right) b \\ & + \left((1-Q) \frac{1+e_A}{2} + Q \frac{1-e_A}{2} \right) QV - \frac{e_A^2}{2} - \frac{e_B^2}{2}. \end{aligned} \quad (31)$$

First-order conditions with respect to e_A and e_B yield,

$$0 = (1-Q)QV + \frac{b(2Q-1)}{2} - e_A, \quad (32)$$

$$0 = -e_B, \quad (33)$$

yielding the solution $\{e_A, e_B\} = \left\{ (1-Q)QV + \frac{b(2Q-1)}{2}, 0 \right\}$. It may be quickly verified that the condition in (29) is satisfied for a sufficiently small b ,

$$E(\theta_B|s_B = L, e_B) - E(\theta_A|s_A = L, e_A) = \left(Q - \frac{Q \frac{1-e_A}{2}}{Q \frac{1-e_A}{2} + (1-Q) \frac{1+e_A}{2}} \right) V \quad (34)$$

$$= \frac{Q(1-Q)V}{\frac{1}{2e_A} - (Q - \frac{1}{2})} - b, \quad (35)$$

which is strictly positive if, and only if,¹¹

$$b < b' \equiv \frac{1 - 2(1-Q)Q(2Q-1)V + \sqrt{1 - 4(1-Q)Q(2Q-1)V}}{(2Q-1)^2}. \quad (36)$$

¹¹Technically, there are two solutions for b' . However, one solution of b' can be ruled out if $e_A > 0$ (which only occurs if $Q \geq \frac{1}{2}$ as is shown shortly), because it is weakly negative valued.

The likelihood of choosing candidate A with $\{e_A, e_B\} = \left\{(1-Q)QV + \frac{b(1-2Q)}{2}, 0\right\}$ is,

$$\Pr(s_A = H|e_A) = Q\frac{1+e_A}{2} + (1-Q)\frac{1-e_A}{2} \quad (37)$$

$$= Q\frac{1+(1-Q)QV + \frac{b(2Q-1)}{2}}{2} + (1-Q)\frac{1-(1-Q)QV - \frac{b(2Q-1)}{2}}{2} \quad (38)$$

$$= \frac{1}{2} + \frac{(2Q-1)(b(2Q-1) + 2(1-Q)QV)}{4}, \quad (39)$$

which is strictly greater than $\frac{1}{2}$ if $Q > \frac{1}{2}$, equals $\frac{1}{2}$ at $Q = \frac{1}{2}$, and is smaller than $\frac{1}{2}$ if $Q < \frac{1}{2}$. The likelihood of choosing candidate B is,

$$\Pr(s_A = L|e_A) = \frac{1}{2} - \frac{(2Q-1)(b(2Q-1) + 2(1-Q)QV)}{4}.$$

Furthermore, the agent's expected utility is given by,

$$E(U) = QV + \frac{b}{2} + \frac{\left((1-Q)QV + \frac{(2Q-1)b}{2}\right)^2}{2}, \quad (40)$$

and the limit of $E(U)$ as $b \rightarrow 0$ is given by,

$$E(U) = QV + \frac{(1-Q)^2 Q^2 V^2}{2}. \quad (41)$$

We have not yet considered the possibility that both the conditions in (17) and (29) do not hold. We now consider that those conditions do not hold: Consider when the conditions in (17) and (29) do not hold and consider the case in which $s_A = L$ and $s_B = H$. The agent's optimal choice is candidate B if, and only if, the condition,

$$E(\theta_B|s_B = H, e_B) - E(\theta_A|s_A = L, e_A) > b, \quad (42)$$

holds. If the conditions in (17) and (29) do not hold and the condition in (42) does hold, then the agent's optimal choice is candidate A if $s_A = s_B = H$, is candidate A if $s_A = H$ and $s_B = L$, and is candidate A if $s_A = s_B = L$. Thus, if $E(\theta_B|s_B = H, e_B) - E(\theta_A|s_A = H, e_A) \leq b$, $E(\theta_B|s_B = L, e_B) - E(\theta_A|s_A = L, e_A) \leq b$, and $E(\theta_B|s_B = H, e_B) - E(\theta_A|s_A = L, e_A) > b$ the

agent's problem is given by,

$$\begin{aligned} \max_{e_A, e_B} & \Pr(s_A = H|e_A) (E(\theta_A|s_A = H, e_A) + b) \\ & + \Pr(s_A = L|e_A) \left(\Pr(s_B = H|e_B) E(\theta_B|s_B = H, e_B) \right. \\ & \left. + \Pr(s_B = L|e_B) (E(\theta_A|s_A = L, e_A) + b) \right) - \frac{e_A^2}{2} - \frac{e_B^2}{2}, \end{aligned} \quad (43)$$

which may be rewritten as,

$$\begin{aligned} \max_{e_A, e_B} & \frac{1+e_A}{2} QV + \left(Q \frac{1+e_A}{2} + (1-Q) \frac{1-e_A}{2} \right) b \\ & + \left((1-Q) \frac{1+e_A}{2} + Q \frac{1-e_A}{2} \right) \left(QV + \left((1-Q) \frac{1+e_B}{2} + Q \frac{1-e_B}{2} \right) b \right) \\ & - \frac{e_A^2}{2} - \frac{e_B^2}{2}. \end{aligned} \quad (44)$$

First-order conditions with respect to e_A and e_B yield,

$$0 = \frac{2(1-Q)QV}{4-b(2Q-1)^2} + \frac{b(2Q-1)}{4+b(2Q-1)^2} - e_A, \quad (45)$$

$$0 = \frac{2(1-Q)QV}{4-b(2Q-1)^2} - \frac{b(2Q-1)}{4+b(2Q-1)^2} - e_B, \quad (46)$$

yielding the solution $\{e_A, e_B\} = \left\{ \frac{2(1-Q)QV}{4-b(2Q-1)^2} + \frac{b(2Q-1)}{4+b(2Q-1)^2}, \frac{2(1-Q)QV}{4-b(2Q-1)^2} - \frac{b(2Q-1)}{4+b(2Q-1)^2} \right\}$. The agent's expected utility is given by,

$$E(U) = QV + b - \frac{b}{4+b(2Q-1)^2} + \frac{(1-Q)^2 Q^2 V^2}{4-b(2Q-1)^2}, \quad (47)$$

and the limit of $E(U)$ as $b \rightarrow 0$ is given by,

$$E(U) = QV + \frac{(1-Q)^2 Q^2 V^2}{4}, \quad (48)$$

which is strictly less than the agent's utilities, as $b \rightarrow 0$, in (28) and (41). Because the agent's expected utilities in (28), (41), and (47) (and her effort choices) are continuous in b , there exists some non-empty range of $b \in [0, \bar{b})$ with $\bar{b} \in (0, \min\{b^*, b'\})$ for which the agent's utility in either

(27) or (40) are strictly greater.¹² For $b \in [0, \bar{b})$, the agent's utility in (27) is greater than her utility in (40) if $Q < 1/2$, the two are equal if $Q = 1/2$ and the agent's utility in (27) is smaller than her utility in (40) if $Q > 1/2$. \square

Proof of Corollary 3. Consider the case with $\{e_A, e_B\} = \{(1 - Q)QV, 0\}$. From Corollary 2, the actor selects item A if and only if $s_A = H$. Thus,

$$\Pr(y = A) = \Pr(s_A = H | e_A = (1 - Q)QV) = \frac{1}{2} + \frac{(2Q - 1)(1 - Q)QV}{2} \quad (49)$$

It follows immediately that $\Pr(y = A) > 1/2$ if $Q > 1/2$ and $\Pr(y = A) < 1/2$ if $Q < 1/2$, and since $e_A > e_B$, this completes the proof. The case with $e_B > e_A$ follows in the same way. \square

Proof of Corollary 4. Differentiating equation (3) with respect to V gives,

$$\frac{d}{dV} \Pr(y = A) = \frac{|2Q - 1|(1 - Q)Q}{2}, \quad (50)$$

which is strictly positive for $Q \neq 1/2$. \square

Proof of Proposition 5. Conditional on exerting effort toward learning about candidate i , the probabilities of i being type $\theta_i = V$, conditional on i being chosen (signal $s_i = H$) and not chosen (signal $s_i = L$) are respectively given by,

$$\Pr(\theta_i = V | y = i) = \frac{\left(\frac{1+e_i}{2}\right) Q}{\left(\frac{1+e_i}{2}\right) Q + \left(\frac{1-e_i}{2}\right) (1 - Q)} \Bigg|_{e_i > 0} \quad (51)$$

$$> Q. \quad (52)$$

$$\Pr(\theta_i = V | y = -i) = \frac{\left(\frac{1-e_i}{2}\right) Q}{\left(\frac{1-e_i}{2}\right) Q + \left(\frac{1+e_i}{2}\right) (1 - Q)} \Bigg|_{e_i > 0} \quad (53)$$

$$< Q. \quad (54)$$

The probabilities of $-i$ being type $\theta_{-i} = V$, conditional on i being chosen (signal $s_i = H$) and not

¹²Note, there exists a sufficiently large b such that the conditions in (17), (29), and (42) do not hold and the solution is trivial: the agent's optimal choice is candidate A always and her optimal effort choices are $\{e_A, e_B\} = \{0, 0\}$, yielding an expected utility of $QV + b$.

chosen (signal $s_i = L$) are respectively given by,

$$\Pr(\theta_{-i} = V \mid y = i) = Q. \quad (55)$$

$$\Pr(\theta_{-i} = V \mid y = -i) = Q. \quad (56)$$

Therefore,

$$\Pr(\theta_i = V \mid y = i) > \Pr(\theta_{-i} = V \mid y = i) = \Pr(\theta_{-i} = V \mid y = -i) > \Pr(\theta_i = V \mid y = -i). \quad (57)$$

To complete the proof, we use the results of Proposition 4 to show the identity of candidate $i \in \{A, B\}$ as determined by Q . □