A BIAS OF SCREENING

David Lagziel and Ehud Lehrer

Discussion Paper No. 19-05

January 2019

Monaster Center for
Economic Research
Ben-Gurion University of the Negev
P.O. Box 653
Beer Sheva, Israel

Fax: 972-8-6472941
Tel: 972-8-6472286

A Bias of Screening*

David Lagziel† and Ehud Lehrer‡

January 17, 2019

Abstract:

This paper deals with the issue of screening. It focuses on a decision maker who, based on noisy unbiased assessments, screens elements from a general set. Our analysis shows that stricter screening not only reduces the number of accepted elements, but possibly reduces their average expected value. We provide a characterization for optimal threshold strategies for screening, and also derive implications to cases where such screening strategies are suboptimal. We further provide various applications of our results to credit ratings, auctions, general trade, the Peter Principle, and affirmative action.

Journal of Economic Literature classification numbers: C70, D49, D81.

Keywords: screening problem; screening bias; threshold strategies.

*The authors wish to thank Tomer Blumkin, David Gilat, Mark Gradstein, Eilon Solan, Rani Spiegler, Ro'i Zultan, two anonymous referees and the associate editor for their valuable comments, as well as the participants of the University of Valencia Economics Seminar, the CORE Econometrics Seminar at UCL, ORSIS 2018, and the Tel-Aviv University Game Theory and Mathematical Economics Research Seminar. Lehrer acknowledges the support of the Israel Science Foundation, Grant #963/15.

†Ben-Gurion University of the Negev, Beer-Sheva 8410501, Israel. e-mail: Davidlag@bgu.ac.il.

‡Tel Aviv University, Tel Aviv 6997801, Israel and INSEAD, Bd. de Constance, 77305 Fontainebleau Cedex, France. e-mail: Lehrer@post.tau.ac.il.

1 Introduction

This paper deals with the problem of screening and, in particular, with the issue of *screening biases*. Our screening problem arises when a decision maker (DM) who filters elements based on noisy unbiased assessments, tries to maximize the average expected value of accepted elements. For that purpose, she fixes a threshold level, namely an acceptance criterion, to filter out low-value elements.

Our first main result concerns the influence of stricter screening on performance. By and large, we show that stricter screening not only reduces the number of accepted elements, but also could lower their expected average value. In other words, a higher bar carries no quality assurances, as a lower one may produce a win-win situation on both sides of the quality vs. quantity (alleged) trade-off.

To exemplify this insight, we use the well-known setting of peer-reviewed academic publishing. Consider a set of academic papers whose potential impact on a standard 12-point grading scale is distributed as in Figure 1. The values and distribution are unknown to the editor and therefore each paper is evaluated via an unbiased noisy refereeing process. The evaluation process generates a mean-preserving spread of the original valuation. Referees evaluate each paper accurately with probability (w.p.) 0.8; otherwise they deviate by two levels, either upwards or downwards, w.p. 0.1 each. As such, the noise is well-defined (given the original distribution of values), independent, symmetric, and of a discrete normal-like distribution. The evaluation distribution is given in Figure 2.

Now, the editor is confronted with the problem of fixing the bar. If she chooses to publish only the top 5%, namely all papers ranked A and above, then the expected value of the published work would be close to $B+$. However, if she lowers the bar to publish the top 13%, roughly all papers ranked $A-$ and above, the inclusive addition of accurately evaluated $A-$ papers would increase the expected value of the published work to $A-$. In that case, not only are additional papers actually published, but their *average* objective impact is higher!

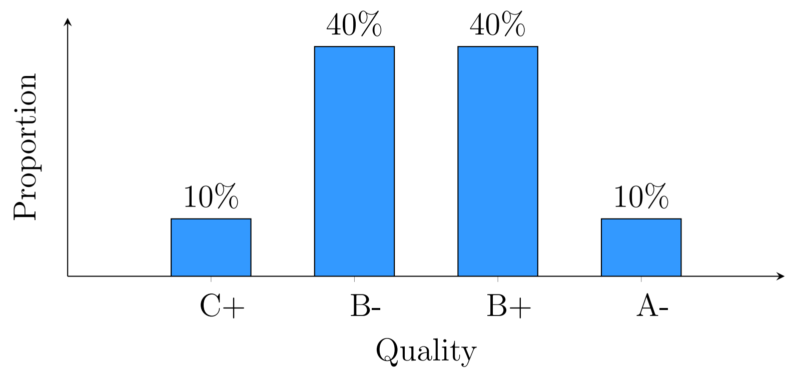

Figure 1: The distribution of papers' impact.

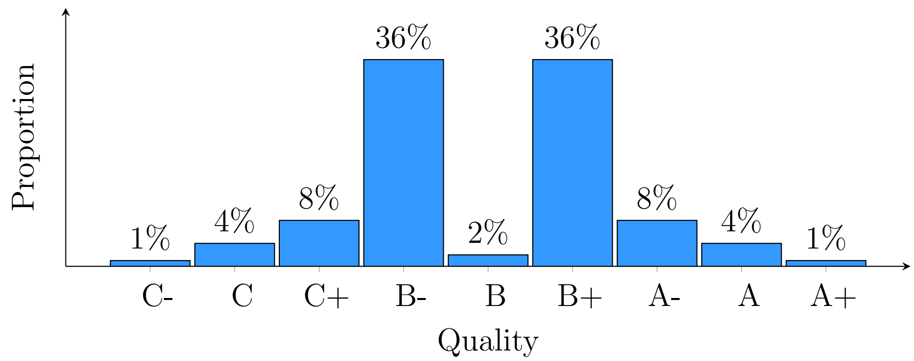

Figure 2: Noisy evaluation of papers' impact. The distribution observed by the editor.

The driving force behind this result is the influence of unbiased noise over different values. A mass of average elements subjected to unbiased noise evidently produces a nominally small, yet relatively significant, amount of upward shifting, whereas a similar effect over a small group of superior elements is mild. To put it differently, *the effect of unbiased noise is potentially biased by the action of screening.* Note that this effect is locally generated and would still hold even if additional papers whose ranks are A and $A+$ are added to the original distribution and their evaluation is completely undistorted.

The above example and insight are based on the key assumption that a threshold strategy has been implemented. Though such strategies are commonly used both in theory and in practice, their optimality has yet to be established. In the example above, a screening strategy that accepts valuations ranked $A-$ and $A+$ would produce a higher average expected value than the threshold strategy which accepts every valuation of at least $A-$. Thus, in the second part of the paper we extend the decision problem to include general utility functions, and address the issue of optimal screening strategies.

The second main result of this paper provides a characterization of optimal threshold strategies for screening. We relate optimal threshold strategies to first- and second-order stochastic dominance of the original valuation, conditional on the noisy assessment at different levels. We then apply our characterization to previous results regarding the Peter Principle (see Peter and Hull (1969)), and derive policy implications concerning affirmative action. Our characterization shows that eliminating affirmative action could lead to a suboptimal outcome.

1.1 Related literature and main contribution

The problem of screening through the examination of pooling and separating equilibria has been well known to economic researchers for decades. The study closest to our own is the seminal work of Stiglitz and Weiss (1981) which investigates the role of interest rates as a screening device. Similarly to our result, they showed that returns are not necessarily monotone with respect to interest rates. Once rates rise, some "safe" borrowers refrain from taking loans, potentially lowering the bank's expected return. Despite these similarities, there is a significant difference between the two non-monotone outcomes. The result of Stiglitz and Weiss (1981) holds in equilibrium whenever borrowers strategically react to the interest-rate mechanism. We, on the other hand, show that non-monotonicity could arise when borrowers are not better informed. In many respects, we respond to and augment the statement of Stiglitz (1975) that "...economies with imperfect information with respect to qualities of individuals differ in fundamental ways from economies with perfect information" by showing that similar phenomena could be attributed to uncertainty rather than asymmetric information. Furthermore, from a technical point of view our results bear some resemblance to the Simpson Paradox in which changes in proportions distort probabilities and produce counter-intuitive results (see, e.g., Simpson (1951), Blyth (1972), and Wagner (1982)).

From the vast literature that considers threshold strategies for screening, the second part of our work strongly relates to the seminal study of Lazear (2004) which provides a theoretical basis for the Peter Principle and proves that consecutive noisy screening leads to an upwards bias. We continue this line of thinking by questioning the incentives for applying threshold strategies. Our results indicate that the upwards bias could be eliminated once optimal non-threshold policies are applied, underlining the adverse effect of suboptimal threshold strategies.

Before we formalize the problem, we wish to emphasize that the widespread use of threshold strategies is well established both in practice and in theory. Whether it be dynamic inventory problem as in Scarf (1959) or admission criteria for top schools, threshold strategies appear to be a natural tool for screening and carry considerable merit. For example, threshold strategies are simple, easy to implement and transparent, thus less vulnerable to manipulation. So regardless of their optimality, the attributes and popularity of threshold strategies make them an important aspect for analysis, as done in the present work.

1.2 Structure of the paper

The paper is organized as follows. In Section 2 we present the basic screening model. In Section 2.1 we establish the first main result concerning screening biases, and in Section 2.2 we derive implications to credit markets, auctions, and general trade. In Section 3 we extend the basic model to tackle the problem of optimal threshold strategies. In Section 3.1 we provide a characterization for these strategies which, in Section 3.2, is applied to several well-known scenarios. Concluding remarks are given in Section 4.

2 The basic screening model

Consider a set of academic papers whose potential impact is distributed according to a non-constant and bounded random variable V, referred to as an *impact variable*. Since the value of each paper is unknown, every paper goes through an unbiased and noisy refereeing process and is publicly evaluated by $V + N$, where N is an *unbiased* random variable, i.e., it is symmetrically distributed around zero and independent of V. The editor uses the noisy evaluation to perform a screening. She fixes a cap $b \in \mathbb{R}$ such that a paper is filtered out if its evaluation is strictly below b. The editor's goal is to maximize the expected impact of the journal (i.e., the average impact of the accepted papers) which is given by

$$\pi(b) = \mathbb{E}[V|V + N \geqslant b].$$

Note that the editor is required to accept some papers in order for the impact to be well-defined. Therefore, any cap b must be *feasible*, meaning that it should not exceed the maximal possible evaluation such that $\Pr(V + N \geqslant b) > 0$. We denote the support of any random variable X by $\left[\underline{X}, \overline{X}\right]$.

Assuming that some papers must be published, the editor is faced with the general problem of fixing the cap. Namely, would it be optimal to set a higher threshold, while maintaining a minimal required volume of publications? A priori, it appears that an increased cap could only improve the journal's expected impact. A higher benchmark screens out less influential papers, driving the average level of accepted papers upwards at the cost of the aggregated impact (since fewer papers are eventually published). To put it differently using various examples: stricter job interviews should improve average production; higher prerequisites should enhance the average student's level; and a more selective choice of products should advance the franchise and increase average prices.

However, this intuition turns out to be false. We prove that a stricter screening not only produces fewer published papers, but could also lower the average level of published ones. Before we turn to solve this paradox, we first need to accurately define a *screening bias*.

Definition 1. *An impact variable V has a* screening bias *if there exists an unbiased noise variable N, such that $\pi(b)$ is non-monotonic. An impact variable V has an* extreme screening bias *if for every positive $\varepsilon < \overline{V} - \mathbb{E}[V]$, there exist a noise variable N and feasible caps $b_2 > b_1$ (which could depend on ϵ) such that*

$$\pi(b_1) \geqslant \overline{V} - \varepsilon > \mathbb{E}[V] = \pi(b_2).$$

In simple terms, an impact variable has no screening biases if, independently of the noise, the expected impact of the published work cannot decrease as the bar is set higher. These biases are extreme if one can generate an almost optimal screening with b_1, while an increase to b_2 generates a result equivalent to no screening whatsoever.

2.1 The robustness of screening biases

We begin by establishing the existence of screening biases. Theorem 1 below states that *every* impact variable has extreme screening biases. One way to derive an intuition for this result is to think of a noise of significant magnitude which has a small probability of taking effect. Given a threshold level that captures a mass of accurate assessments, the screening is not significantly affected by the noise. However, for a higher threshold which is fixed above these accurate evaluations, low and high values fluctuate independently such that the screening becomes redundant. We emphasize (and prove later on) that this is only one possibility out of many for biases to emerge.

Theorem 1. *Every impact variable has extreme screening biases.*

Proof. Fix $\varepsilon \in \left(0, \overline{V} - \mathbb{E}[V]\right)$. Without loss of generality (up to a linear transformation of the impact variable), we prove the above statement for an impact variable supported on $[0,1]$. To be specific, we assume that $V \in [0,1]$ w.p. 1, with a CDF F such that $F(1) = 1$, and $\Pr(1-\delta \leqslant V \leqslant 1) > 0$ for every $\delta > 0$.

If $\Pr(V = 1) > 0$, take $\delta_1 \in (0,1)$ and consider $b_1 = 1 - \delta_1 < 1 + \delta_1 = b_2$. Define N by

$$N = \begin{cases} \pm(1+\delta_1), & \text{w.p. } \frac{\delta_1}{2}, \\ 0, & \text{w.p. } 1-\delta_1. \end{cases}$$

It follows that $V + N \geqslant b_2 \Leftrightarrow N = 1 + \delta_1$; and so $\pi(b_2) = \mathbb{E}[V]$. However,

$$\begin{aligned} \pi(b_1) &= \frac{\mathbb{E}\left[V \mathbb{1}_{\{V+N \geqslant 1-\delta_1\}}\right]}{\Pr\left(V+N \geqslant 1-\delta_1\right)} \\ &= \frac{\Pr(N=0)\mathbb{E}[V\mathbb{1}_{\{V \geqslant 1-\delta_1\}}] + \Pr(N = 1+\delta_1)\mathbb{E}[V]}{\Pr(N=0)[1-F\left(1-\delta_1\right)] + \Pr(N=1+\delta_1)} \\ &= \frac{(1-\delta_1)\mathbb{E}[V\mathbb{1}_{\{V \geqslant 1-\delta_1\}}] + \frac{\delta_1}{2}\mathbb{E}[V]}{(1-\delta_1)[1-F\left(1-\delta_1\right)] + \frac{\delta_1}{2}} \to \mathbb{E}[V|V \geqslant 1], \quad \text{as} \quad \delta_1 \to 0. \end{aligned}$$

Hence, the result holds true for a sufficiently small $\delta_1 > 0$.

If $\Pr(V = 1) = 0$, follow the same computation with $0 < \delta_1, \delta_2 < 1$, caps $b_1 = 1 - \delta_2 < 1 = b_2$, and

$$N = \begin{cases} \pm 1, & \text{w.p. } \frac{\delta_1}{2}, \\ 0, & \text{w.p. } 1-\delta_1. \end{cases}$$

The computation yields $\pi(b_2) = \mathbb{E}[V]$, and $\lim_{\delta_1 \to 0} \pi(b_1) = E[V|V \geqslant 1-\delta_2]$. By taking sufficiently small δ_1 and δ_2, the statement of Theorem 1 holds. ■

Remark 1. *As a technical generalization, one can partially extend* Theorem 1 *to show that all random variables, including unbounded ones, have screening biases.*

Once existence of screening biases is established, we move on to the important issue of robustness. In general, one could argue that the phenomenon of screening biases is restricted to significant noises that completely distort the screening. We approach this concern in several ways. First, note that even mild noises instigate biases. For example, the noise used in the proof of Theorem 1 is not only unbiased, but maintains an infinitesimal variance. Second, in Claims 1 and 2 below we demonstrate the universal nature of the relevant noise variables. In particular, Claim 1 considers a normal-like continuous[1] noise — with a high probability it is uniformly distributed on a small interval around zero; otherwise, it is uniformly distributed on a broader interval around zero. The claim states that for every impact variable one can find such bias-generating noises.

Claim 1. *Every impact variable has a continuous noise variable that produces screening biases.*

Proof. Similarly to Theorem 1, consider V on $[0,1]$ with a CDF F where $\Pr(1-\delta \leqslant V \leqslant 1) > 0$ for every $\delta > 0$, and $\Pr(V \in (0,1)) > 0$. Fix $\varepsilon, c \in (0,1)$ such that N is, w.p. c, distributed uniformly on $(-\varepsilon, \varepsilon)$, and w.p. $1-c$ is distributed uniformly on $(-1,-\varepsilon) \cup (\varepsilon, 1)$. Take $b_1 = 1$ and $b_2 = 1+\varepsilon$. A direct computation shows that

$$\begin{aligned}
\pi(b_2) &= \frac{\int_{\varepsilon}^{1} v(v-\varepsilon)dF(v)}{\int_{\varepsilon}^{1}(v-\varepsilon)dF(v)}, \\
\pi(b_1) &= \frac{\frac{1-c}{2-2\varepsilon}\int_0^{1-\varepsilon} v^2 dF(v) + \frac{1}{2\varepsilon}\int_{1-\varepsilon}^{1} v[\varepsilon + c(v-1)]dF(v)}{\frac{1-c}{2-2\varepsilon}\int_0^{1-\varepsilon} v dF(v) + \frac{1}{2\varepsilon}\int_{1-\varepsilon}^{1}[\varepsilon + c(v-1)]dF(v)}.
\end{aligned}$$

If $c \to 1$, then

$$\pi(b_1) \to \frac{\int_{1-\varepsilon}^{1} v[\varepsilon + v - 1]dF(v)}{\int_{1-\varepsilon}^{1}[\varepsilon + v - 1]dF(v)}.$$

In addition, if $\varepsilon \to 0$, then $\pi(b_1) \to 1$, while $\pi(b_2) < 1$. Thus, Claim 1 holds for ε and c sufficiently close to 0 and 1, respectively. ∎

Claim 1 shows that the DM can have a very accurate screening system, such that errors are mild in terms of magnitude and plausibility and still encounter biases. Moreover, the fact that the screening process is highly accurate (but still noisy) actually enables biases to emerge. The expected value is sensitive to changes in valuations relative to their mass; therefore biases are amplified by the fact that valuations are highly accurate although some errors are still possible.

The next result further extends the last conclusion. By restricting the discussion to finitely discrete distributions, one can prove that biases exist as long as the noise is not completely negligible relative to the impact variable. Formally, given finitely supported variables V and N, the noise variable is considered V-*distinguishable* if for every two adjacent values, $v_1 < v_2$, of the impact variable it follows

[1] We use the term "continuous" to describe a non-atomic random variable that is fully supported on an interval.

that $|\overline{N} - \underline{N}| > v_2 - v_1$. That is, a noise variable is V-distinguishable if its support cannot be bounded by two adjacent values of V. The following claim shows that any distinguishable noise produces biases.

Claim 2. *For every finitely supported impact and noise variables where the noise is V-distinguishable, there exists $\alpha \in (0,1)$ such that αN produces screening biases.*

Proof. Consider V and N where N is V-distinguishable, $\text{Supp}(V) = \{v_1 < v_2 < \cdots < v_m\}$, and $\text{Supp}(N) = \{n_1 < n_2 < \cdots < n_k\}$. Denote $\beta_{ij} = v_i + n_j$ for any i and j. Since the two variables have finite support, we can take any contraction αN such that the β_{ij}s are distinct. Fix a cap of $b_1 = \beta_{m1}$. Then

$$\Pr(V + N \geqslant b_1) = \Pr(V = v_m,\ N = n_1) + \sum_{n > n_1} \Pr(V > \beta_{m1} - n, N = n),$$

so that the second term in the RHS contains positive-probability events where $V < v_m$ due to the assumption on N.

If the cap is infinitesimally increased to $b_2 > b_1$, without crossing the following β_{ij}, we get $\Pr(V + N \geqslant b_2) = \Pr(V + N \geqslant b_1) - \Pr(V = v_m,\ N = n_1)$, and the probability $\Pr(V = v_m | V + N \geqslant b)$ decreases. Since $v_m = \max\{V\}$, the conditional probability of non-maximal values of V increases as well. Thus, the weighted average of V (given $V + N \geqslant b$) decreases if the cap is increased from b_1 to b_2, and Claim 2 follows. ■

The role of the scaling factor α is twofold. On the one hand, it is technically used to eliminate potential equalities and ensure that all assessments are distinct. On the other hand, the factor shows that a high-magnitude noise is not a necessary condition for a screening bias. These two attributes exemplify how the noise magnitude is relevant only to the extent to which it enables strict ordinal changes of the original valuations. Moreover, the proof of Claim 2 is based on the top two values of V (i.e., $v_2 = \max\{V\}$), so one can weaken the condition on the distinguishable noise to these two values specifically.

It is not a mere coincidence that the top valuations are essential for the proof of Claim 2. The result of the following lemma builds on this notion, showing that biases are likely to emerge at a top-level screening rather than at a lower one. Specifically, biases do not occur when thresholds are bounded from above by $\mathbb{E}[V] + \underline{N}$, which implies biases are more prevalent in screening at high levels. In addition, Lemma 1 shows that the magnitudes of biases are bounded by the size of the noise, meaning that the decrease in expected impact is limited to the support of the noise.

Lemma 1. *If $b < \mathbb{E}[V] + \underline{N}$, then the expected impact $\pi(b)$ is an increasing function. Moreover, for every two feasible caps $b_2 > b_1$, we have that $\pi(b_2) \geqslant \pi(b_1) - |\overline{N} - \underline{N}|$.*

Proof. Fix V, N, and $b_1 < b_2 < \mathbb{E}[V] + \underline{N}$. A necessary and sufficient condition for a bias is $\mathbb{E}[V | V + N \in [b_1, b_2)] > \mathbb{E}[V | V + N \geqslant b_2]$. Since the LHS equals $\mathbb{E}[V | b_1 - N \leqslant V < b_2 - N]$ and

$b_2 - N < \mathbb{E}[V]$ w.p. 1, we conclude that the LHS is bounded by $\mathbb{E}[V]$ while $\mathbb{E}[V|V+N \geqslant b_2] \geqslant \mathbb{E}[V]$ (as needed for Lemma 1).

For the second part of the lemma, fix feasible caps $b_2 > b_1$, and define $Y = V + N$. Then

$$\begin{aligned}
\pi(b_2) - \pi(b_1) + \overline{N} - \underline{N} &= \mathbb{E}[Y - N|Y \geqslant b_2] - \mathbb{E}[Y - N|Y \geqslant b_1] + \overline{N} - \underline{N} \\
&\geqslant \mathbb{E}[N|Y \geqslant b_1] - \mathbb{E}[N|Y \geqslant b_2] + \overline{N} - \underline{N} \\
&\geqslant \underline{N} - \overline{N} + \overline{N} - \underline{N} > 0,
\end{aligned}$$

which concludes the proof. ■

Both statements follow the same motivation as Claim 2 — the ordinal changes are bounded by the magnitude of the noisy deviation, and filtering at bottom levels cannot generate a significant impact. The leading conclusion, therefore, is that designers should not concern themselves with biases whenever screening is restricted to low-level elements and when the potential loss from a limited noise is negligible.

2.1.1 Screening biases and the monotone likelihood ratio property

Two probability distributions are said to have the monotone likelihood ratio property (MLRP) if their ratio is non-decreasing. In our context, one can take two realized signals $s_1 > s_2$ and consider the ratio between the distributions of $V|\{V + N = s_1\}$ and $V|\{V + N = s_2\}$. Under the MLRP, this ratio is a non-decreasing function. Moreover, if the MLRP holds, then the distribution of $V|\{V + N = s_1\}$ first-order stochastically dominates the distribution of $V|\{V + N = s_2\}$, which leads to a higher expected value (i.e., $\mathbb{E}[V|V + N = s_1] > \mathbb{E}[V|V + N = s_2]$). Therefore, the law of total expectation implies that a screening bias can exist only if the MLRP fails. This observation becomes useful in Section 3, where we provide a weaker property than MLRP to eliminate screening biases, which also characterizes optimal threshold strategies.

2.2 Implications of screening biases

2.2.1 Credit ratings

Potential errors in credit ratings occur more frequently than thought. In many cases, either high- or low-risk debtors are wrongfully flagged as low- and high-risk ones, thus posing a substantial problem in credit markets.[2] To address this problem using our formulation, assume that V defines the actual solvency of a group of debtors and N denotes the evaluation error of the rating agency. Thus, screening biases correspond to stricter loan conditions that generate lower expected returns. Typically, this is

[2] For example, see the *CNBC* article, "The real problem with credit reports is the astounding number of errors", Aaron Klein, September 27, 2017.

not surprising in credit markets. Stricter screening regularly limits the potential profit by reducing the risk. Yet, our screening biases are fundamentally different since the relation between returns and risk via the interest-rate mechanism is irrelevant. In our context, the lower expected returns are due to bad loans, rather than safer ones. In other words, a stricter screening eliminates high-value debtors who could increase expected profits, while reducing the creditor's risk.

Building on the result of Theorem 1, extreme screening biases suggest that not only expected returns decrease, but the investment becomes riskier. Specifically, once $\pi(b)$ is ε-close to optimality, the variance over accepted valuations is close to zero. This result follows directly from the Bhatia-Davis inequality (see Bhatia and Davis (2000)), stating that $\mathrm{Var}(X) \leqslant \left(\overline{X} - \mathbb{E}[X]\right)\left(\mathbb{E}[X] - \underline{X}\right)$ for every bounded random variable X. Therefore, for every extreme screening bias with a sufficiently small ε, the transition from a low cap to a higher one increases the variance over the investment and decreases the expected return: a lose-lose situation.

2.2.2 Auctions and the winner's curse

Aside from the general screening problem, our model also applies to auctions. For example, consider an English auction where b is the starting price. Under uncertainty regarding the value of the auctioned item, whether it is a given commodity in regular auctions or a certain project in procurement auctions, the opening price is of key importance. Assuming that the valuation among bidders is given by V, and bidders' uncertainty is projected through N, then the opening bid determines the distribution of valuations among the auction participants. Only bidders with noisy valuations above the cap (i.e., $V + N \geqslant b$) would participate, and screening biases suggest that a higher opening bid is not necessarily productive and does not guarantee efficiency.

The implication of screening biases to auctions is related to the well-known phenomenon of the winner's curse. The fact that bids depend on a mean-preserving spread of the true valuation suggests that the winning bidder may not be the one with the highest valuation, but only the most optimistic one. Such inefficiency becomes crucial in credit auctions,[3] as a higher cap possibly increases the probability of a default.

2.2.3 General trade

The last implication we provide relates to general trade. Let V represent the intrinsic quality of products in a given market, whereas $V + N$ reflects the subjective assessments among sellers. If b is the current price, then only sellers with a valuation of $V + N \leqslant b$ will agree to sell. With a few algebraic adjustments, it follows that a screening bias is equivalent to non-monotonicity of $\mathbb{E}[V|V + N \leqslant b]$

[3] Credit auctions are auctions where payments are not immediate, as in long-term infrastructure and real-estate projects. See Parlane (2003), Board (2007), and Lagziel (2018), among many others.

with respect to the price b. Namely, a higher price may reduce the expected quality of products in the market by introducing relatively more low-value products with overzealous sellers.

3 The optimality of threshold strategies

The model and results of Section 2 are based on two key assumptions: (i) the DM's goal is to maximize the expected value of V; and (ii) the DM uses a threshold strategy. Though both assumptions seem straightforward, one should bear in mind that the DM can exercise more sophisticated policies to maximize a utility function that depends on V. In the current section, we extend the basic model by considering an expected-utility maximizer DM who faces a screening problem subject to the information encapsulated in $V + N$. Our goal is to establish conditions under which threshold strategies are indeed optimal, and study the practical implications in cases where they are suboptimal. Doing so, we also establish a relation between the optimality of threshold strategies and the preceding analysis of screening biases.

3.1 The extended model and main result

Fix an impact variable V and a noise variable N, and consider a DM with a utility function $u : \mathbb{R} \to \mathbb{R}$. As an expected-utility maximizer, the DM sets a screening strategy $\sigma : \mathbb{R} \to \{0, 1\}$, where 1 denotes the acceptance of a specific valuation and 0 denotes a rejection. Given σ, the DM's expected utility is

$$u(\sigma) = \mathbb{E}\left[u(V)\mathbb{1}_{\{\sigma(V+N)=1\}}\right].$$

That is, the DM tries to maximize the expected utility from the accepted elements, subject to the noisy evaluation $V+N$. Evidently, the decision to accept an element hinges on the sign of the utility function, thus the DM would accept any positive-probability set A of valuations if $\mathbb{E}\left[u(V)\mathbf{1}_{\{V+N\in A_0\}}\right] \geqslant 0$ for every subset $A_0 \subseteq A$. Note that all definitions and statements hold almost surely, i.e., hold up to a zero-measure deviation.

A strategy σ is a *threshold* (cut-off) strategy if the acceptance of some values with positive probability implies that higher valuations are not rejected with a positive probability. Formally, a strategy σ is a threshold strategy if for every t, the condition $\Pr(s \leqslant t : \ \sigma(s) = 1) > 0$ implies that $\Pr(s > t : \ \sigma(s) = 0) = 0$. Define σ_u to be the optimal strategy given u, where the notion of optimality is taken in the usual sense that $u(\sigma_u) \geqslant u(\sigma)$ for every strategy σ.

We can now address the question of optimal threshold strategies and, specifically, to characterize the conditions ensuring that such strategies are optimal. For that purpose, the following notations are needed. For every $A \subseteq \mathbb{R}$, let $V|_A$ be the conditional distribution of V given $V + N \in A$. Denote first-order (second-order) stochastic-dominance by $\succeq_{\mathrm{I}}$ (respectively, $\succeq_{\mathrm{II}}$). The following theorem provides a characterization for optimal threshold strategies, based on the stochastic dominance of V conditional on different feasible valuations.

Theorem 2. *The optimal strategy σ_u is a threshold strategy for every increasing (increasing and concave) utility function u if and only if for every two positive-probability sets $A, B \subseteq \mathbb{R}$ such that $a > b$ for every $a \in A$ and $b \in B$, it holds that $V|_A \succeq_{\mathrm{I}} V|_B$ (respectively, $V|_A \succeq_{\mathrm{II}} V|_B$).*

Proof. We prove the first part of the theorem, and the second part follows similarly by taking an increasing and concave utility function. For every two sets $A, B \subseteq \mathbb{R}$, denote $A > B$ if $a > b$ for every $a \in A$ and for every $b \in B$.

For the first part we use a proof by contradiction. Assume that for every two sets $A > B$ of positive probability we have that $V|_A \succeq_{\mathrm{I}} V|_B$, and that there exists an increasing utility function u such that σ_u is not a threshold strategy. Namely, there exists t_0 such that $\Pr(s \leqslant t_0 : \ \sigma(s) = 1) > 0$ and $\Pr(s > t_0 : \ \sigma(s) = 0) > 0$. Let $A = \{s > t_0 : \ \sigma(s) = 0\}$ and $B = \{s \leqslant t_0 : \ \sigma(s) = 1\}$. The sets A and B are of positive probability and, by definition, $A > B$. FOSD implies that $\mathbb{E}[u(V)|V + N \in A] \geqslant \mathbb{E}[u(V)|V + N \in B]$. However, $\sigma(A) = 0$ and $\sigma(B) = 1$ suggest that $\mathbb{E}[u(V)\mathbf{1}_{\{V+N\in B\}}] \geqslant 0 > \mathbb{E}[u(V)\mathbf{1}_{\{V+N\in A\}}]$, which contradicts the previous inequality.

For the converse, assume there exist sets $A > B$ of positive probability such that $V|_A \not\succeq_{\mathrm{I}} V|_B$. Thus, there exists an increasing utility function u such that $\mathbb{E}[u(V)|V + N \in B] > \mathbb{E}[u(V)|V + N \in A]$. Consider an auxiliary utility function $\tilde{u} = u - c$ where $\mathbb{E}[\tilde{u}(V)|V + N \in B] > 0 > \mathbb{E}[\tilde{u}(V)|V + N \in A]$, which implies that $\mathbb{E}[\tilde{u}(V)\mathbf{1}_{\{V+N\in B\}}] > 0 > \mathbb{E}[\tilde{u}(V)\mathbf{1}_{\{V+N\in A\}}]$. Hence, $\sigma_{\tilde{u}}(A) = 0$ and $\sigma_{\tilde{u}}(B) = 1$, up to a zero-measure deviation. Therefore, $\sigma_{\tilde{u}}$ is not a threshold strategy. ■

Let us explain Theorem 2 in simple terms. Consider any two sets of evaluations A and B both of positive probability such that A is (point-wise) above B. Assuming that $V|_A$ stochastically dominates $V|_B$, then threshold strategies are optimal since a non-negative expected payoff at the lower values leads to a non-negative expected payoff at higher ones. As it turns out, the other direction holds just as well: once stochastic dominance fails, one can easily construct a relevant utility function which does not obey the threshold-strategy optimality criteria.

Before we elaborate on the applicable aspects of Theorem 2, we relate it to the results given in Section 2. In particular, note that the necessary condition for optimal threshold strategies given in Theorem 2 disallows screening biases to emerge. That is, if a screening bias exists, then there exist sets $A > B$ of positive probability such that $\mathbb{E}[V|_B] > \mathbb{E}[V|_A]$. The last condition is weaker than MLRP (as discussed in Section 2.1.1), since monotonicity is required only in terms of expected values rather than in terms of the ratio of distributions.

The combination of Theorems 1 and 2 raises the question of optimal threshold strategies that still allow screening biases to emerge. The answer to this question is based on the chosen utility function. Since a screening bias contradicts the condition given in Theorem 2, one can construct a utility function such that threshold strategies are suboptimal. This, however, does not imply that threshold strategies should never be used since, for some utility functions, threshold strategies are indeed optimal (for example, for any positive function).

3.2 Applications

3.2.1 The Peter Principle

The Peter Principle was initially coined by Peter and Hull (1969) and later studied in a theoretical framework by Lazear (2004) in the context of consecutive screening. Lazear showed that regression to the mean prompts an upwards bias once a threshold strategy is used. This bias explains numerous phenomena, ranging from failed managerial promotions to unsuccessful movie sequels. The threshold strategies in Lazear's model ensure that accepted elements have a temporal advantage over rejected ones due to the realized noise. However, such strategies are potentially suboptimal, whereas the implementation of optimal non-threshold policies can completely eliminate the mentioned bias. Let us consider a concrete example to explain this phenomenon.

Take a uniformly distributed impact variable $V \sim U(-0.5, 0.5)$, and a symmetric binary noise variable $N = \pm 0.5$. For simplicity, assume that $u(x) = x$. Fix a threshold strategy $\sigma = \mathbf{1}_{[0,1]}$ which screens out every negative valuation. This strategy would produce the same uniform distribution as the original one, i.e., $V|_{[0,1]} \sim U(-0.5, 0.5)$. As Lazear points out, the expected noise of every accepted (rejected) element is positive (negative, respectively), so the upward bias is evident. Formally, the computation shows that

$$\mathbb{E}[N|\sigma(v+N) = 1] = 0.5 < -0.5 = \mathbb{E}[N|\sigma(v+N) = 0].$$

However, this analysis follows a threshold strategy that is clearly suboptimal.

Consider the strategy $\sigma = \mathbf{1}_{[-0.5,0] \cup [0.5,1]}$ which accepts elements of noisy valuation between $[-0.5, 0]$ and $[0.5, 1]$. This strategy yields an expected noise $\mathbb{E}[N|\sigma(V + N) = 1] = 0$ and a positive expected value $\mathbb{E}[V\mathbf{1}_{\{\sigma(V+N)=1\}}] = \frac{1}{8}$. In other words, a non-threshold strategy produces positive expected value with no expected bias. It is important to note that the latter strategy is indeed optimal, and strictly dominates any threshold strategy. The same non-monotonicity holds for other limited noises such as $N = \pm n$ where $n > 0.25$. This simple example illustrates how suboptimal threshold strategies generate a bias that optimal non-threshold policies can overcome.

3.2.2 Affirmative action

Affirmative action advocates the promotion of education and employment for discriminated individuals of certain groups. These policies are aimed at levelling the playing field, i.e., maintaining equal opportunities. To illustrate this notion in our model, consider a continuous impact variable V which denotes the true valuation of individuals in a certain position, and a finitely supported N which depends on various irrelevant characteristics. In our formulation, affirmative action is manifested through different screening criteria among heterogeneous individuals, namely the screening is based on a non-threshold strategy, while the absence of affirmative action suggests a threshold strategy.

To exemplify this idea, consider the previous example where $V \sim U(-0.5, 0.5)$ and $N = \pm 1/3$ with equal probabilities. A straightforward computation shows that the optimal screening strategy is $\sigma = \mathbf{1}_{[0,1/6] \cup [1/3,5/6]}$ and not a threshold strategy. The $\left[0, \frac{1}{6}\right]$-acceptance condition targets negative noise realizations, while the $\left[\frac{1}{3}, \frac{5}{6}\right]$-condition targets positive ones. This optimal strategy specifically *targets individuals with different noise realizations*, illustrating the positive economics of affirmative action which go beyond the normative ones. Claim 3 generalizes this example.

Claim 3. *If V is continuous and N has finite support, there exists positive-probability sets $A > B$ such that $V|_A \not\succeq_{\mathrm{I}} V|_B$.*

Proof. Fix V and N as stated such that $\mathrm{Supp}(N) = \{n_1 < n_2 < \cdots < n_k\}$. If $n_k + \underline{V} \geqslant \overline{V} + n_{k-1}$, then $V|_{[n_k + \underline{V}, n_k + \overline{V}]} \sim V$ which does not dominate $V|_{[\overline{V} + n_{k-1} - \varepsilon, \overline{V} + n_{k-1}]}$, for a sufficiently small $\varepsilon > 0$. Therefore, consider $\delta = n_k - n_{k-1} < \overline{V} - \underline{V}$, and fix $\epsilon = \frac{\delta}{2}$. Denote $\overline{V}_{k-1} = \overline{V} + n_{k-1}$, and define the intervals $A = (\overline{V}_{k-1}, \overline{V}_{k-1} + \epsilon)$ and $B = (\overline{V}_{k-1} - \epsilon, \overline{V}_{k-1})$. By definition $A > B$, but $\mathrm{Supp}(V|_A) = (\overline{V} - \delta, \overline{V} - \varepsilon)$ while $\mathrm{Supp}(V|_B) = (\overline{V} - \delta - \varepsilon, \overline{V} - \delta) \cup (\overline{V} - \varepsilon, \overline{V})$. Hence, $\Pr(V > \overline{V} - \varepsilon | B) > 0 = \Pr(V > \overline{V} - \varepsilon | A)$ and the claim follows. ■

Note that the proof carries some resemblance to Lemma 1 in the sense that the strategies' discontinuity is imminent at top valuations.

4 In conclusion

This paper deals with two screening-related problems: the first concerns the adverse effect of stricter screening and the second focuses on the actual screening method. In the first part of the paper we establish the existence and robustness of screening biases, and in the second part we characterize optimal threshold strategies. We combine the two problems and results by showing that screening biases exist as long as the necessary and sufficient condition for optimal threshold strategies fails. Thus, although the two parts deal with slightly different goals, they lead to similar conditions overall.

References

Bhatia, R. and C. Davis (2000): "A better bound on thevVariance," *The American Mathematical Monthly*, 107, 353.

Blyth, C. R. (1972): "On Simpson's Paradox and the Sure-Thing Principle," *Journal of the American Statistical Association*, 67, 364–366.

Board, S. (2007): "Bidding into the Red: A Model of Post-Auction Bankruptcy," *The Journal of Finance*, 62, 2695–2723.

Lagziel, D. (2018): "Credit auctions and bid caps," *Games and Economic Behavior*.

Lazear, E. P. P. (2004): "The Peter Principle: A theory of decline," *Journal of Political Economy*, 112, S141–S163.

Parlane, S. (2003): "Procurement Contracts under Limited Liability," *Economic and Social Review*, 34, 1–22.

Peter, L. J. and R. Hull (1969): *The Peter Principle: Why Things Always Go Wrong*, William Morrow & Company, New York.

Scarf, H. (1959): "The optimality of (S, s) policies in the dynamic inventory problem," in *Mathematical Models in Social Sciences*, ed. by K. J. Arrow, S. Karlin, and P. Suppes, Stanford University Press.

Simpson, E. H. (1951): "The Interpretation of Interaction in Contingency Tables," *Journal of the Royal Statistical Society. Series B (Methodological)*, 13, 238–241.

Stiglitz, J. E. (1975): "The Theory of "Screening," Education, and the Distribution of Income," *American Economic Review*, 65, 283–300.

Stiglitz, J. E. and A. Weiss (1981): "Credit rationing in markets with imperfect information," *The American Economic Review*, 71, 393–410.

Wagner, C. H. (1982): "Simpson's Paradox in real life," *The American Statistician*, 36, 46–48.