
Management Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

# The Worst-First Heuristic: How Decision Makers Manage Conjunctive Risk

Joshua Lewis, Daniel Feiler, Ron Adner

informs
http://pubsonline.informs.org/journal/mnsc

MANAGEMENT SCIENCE
*Articles in Advance*, pp. 1–22
ISSN 0025-1909 (print), ISSN 1526-5501 (online)

# The Worst-First Heuristic: How Decision Makers Manage Conjunctive Risk

**Joshua Lewis,[a,]* Daniel Feiler,[b] Ron Adner[b]**

[a] New York University, New York, New York 10012; [b] Dartmouth College, Hanover, New Hampshire 03755
*Corresponding author
**Contact:** joshua.lewis@nyu.edu, https://orcid.org/0000-0002-9114-2711 (JL); daniel.feiler@dartmouth.edu, https://orcid.org/0000-0001-6248-5654 (DF); ron.adner@dartmouth.edu (RA)

**Received:** March 20, 2021
**Revised:** October 3, 2021; December 3, 2021
**Accepted:** December 13, 2021
**Published Online in Articles in Advance:** April 6, 2022

**https://doi.org/10.1287/mnsc.2022.4411**



**Abstract.** Many important managerial outcomes hinge on the co-occurrence of multiple uncertain events, a situation termed *conjunctive risk*. Whereas past literature has addressed the psychology of *choosing* to enter situations with conjunctive risk, this article elucidates a novel way in which the psychology of *managing* conjunctive risk is importantly distinct. We examine a case in which there are two independent events, one is currently less likely than the other, both are required for overall success, and the decision maker must evaluate opportunities to increase the chance of the less-likely or more-likely requirement. We introduce the hypothesis of a worst-first heuristic. Decision makers intuitively evaluate improvements in conjunctive risk according to their impact on the biggest barrier to success, the least likely of the required events. We find evidence for such a worst-first heuristic across nine experiments ($n$ = 3,653, including samples from the United States and United Kingdom in Studies 1–5 and Studies S1–S3 in the online supplement, as well as a sample of managers in Study 6). Participants invest more to improve chances of less-likely requirements than more-likely requirements, even when the latter improvements have at least as much impact on the aggregate chance of success. Moreover, we find that decision makers exhibit this behavior particularly when *managing* conjunctive risk, as doing so makes them attend to which threat is the worst. Conversely, they do not appear to exhibit the behavior when making formally equivalent decisions about *choosing* between conjunctive risks. This bias toward underinvesting in stronger-links holds important implications for decision making in contexts subject to conjunctive risk—both managerial and societal.

**History:** Accepted by Yuval Rottenstreich, behavioral economics and decision analysis.
**Supplemental Material:** Data and the e-companion are available at https://doi.org/10.1287/mnsc.2022.4411.

**Keywords:** decision analysis • risk • judgment and decision making • heuristics and biases • conjunctive events

## Introduction

Success is a matter of conjunctive risk when it depends on multiple uncertain events all occurring, as is common, across a wide variety of managerial decision settings. For example, most projects require the successful execution of multiple, distinct components, whose completion is uncertain. A large literature has examined how decision makers evaluate opportunities to take conjunctive risks (Bar-Hillel 1973; Nilsson et al. 2009, 2013; Adner and Feiler 2019) or situations with compound risk more broadly (Budescu and Fischer 2001, Fan et al. 2018). Underexplored, however, is the question of how decision makers *manage* conjunctive risk. How do decision makers invest time, effort, or money to improve their overall chance of success when that success depends on multiple uncertain events? In addressing this question, we make two primary contributions. First, we delineate a novel heuristic that decision makers use to manage conjunctive risk. Second, we highlight how decision makers' choices can differ systematically according to whether they are *managing* conjunctive risks or, as has been the focus of past literature, *choosing* between conjunctive risks. This article offers unique insight into how active managers of conjunctive risk perceive it, evaluate it, and influence it.

We operationalize managing conjunctive risk as a case in which:

1. A decision maker will receive a payoff only if two independent, uncertain requirements are both met (and no payoff if either is *not* met).

2. One requirement is currently less likely than the other (as is typically the case). We will call the chance of the less-likely requirement the *weaker-link* and the chance of the more-likely requirement the *stronger-link*.

3. The decision maker can take an action to improve either the weaker-link or the stronger-link.

This case captures the essence of managing conjunctive risk. The decision makers must determine whether and how to invest resources in improving the weaker-link and the stronger-link to maximize their chance of overall success.

Consider an example in which you require two events to occur to earn a payout: one requirement has a 20% chance (the weaker-link) and the other statistically independent requirement has a 40% chance (the stronger-link). Because you need both independent events to occur, you currently have an 8% chance of succeeding (20% × 40%). Now imagine that you can choose which of two improvements to invest in: either you can improve the weaker-link by 15 percentage points (from 20% to 35%) or the stronger-link by 30 percentage points (from 40% to 70%). Which would you choose?

Rationally, a decision maker should be indifferent between these two improvements. Even though the 15-percentage-point improvement in the weaker-link is only half the size of the 30-percentage-point improvement to the stronger-link, they both constitute the same *proportional* change with respect to the constituent probability they are improving (30/40 = 15/20) and, as a consequence, the same aggregate chance of success (14%).[1] Previous literature suggests that participants might fail to appreciate this subtlety. First, a common theme of decision-making research is that people often process information in the simplest way possible (Slovic 1972), as if "what you see is all there is" (Kahneman 2011, p. 85). Second, decision makers attend to absolute magnitudes even when they should not (e.g., Burson et al. 2009, Shue and Townsend 2021). According to this work, decision makers might mistakenly believe that it was always best to choose the greatest increase in the chance of meeting any individual requirement. Per the example above, because the 30-percentage-point improvement in the stronger-link is of larger magnitude than the 15-percentage-point improvement in the weaker-link, the former might seem better: after all, 30 > 15. By this logic, rather than correctly thinking about these improvements in proportional terms, decision makers will naively attend to their absolute magnitude, thereby predictably overvaluing improvements in stronger-links relative to improvements in weaker-links. Thus, based on this literature, we initially expected such a bias to be the dominant tendency.

Although we do find evidence that *some* participants behaved in accordance with this expectation (to which we will return later in the article), on aggregate, the *opposite* tendency was dominant—overvaluation of improvements in weaker-links. For example, not only do decision makers prefer a 15-percentage-point improvement in a 20% weaker-link over an equally valuable 30-percentage-point improvement in a 40% stronger-link, but they even prefer it over an objectively-more-valuable 34-percentage-point improvement in the stronger-link, which violates axiomatic rational choice.

We propose that this result is the product of the following decision process. Despite recognizing that they cannot simply judge improvements by their absolute magnitude, decision makers are still often unable or unwilling to calculate and rationally assess the implications of potential improvements for their aggregate chance of success. Instead, they form an initial *intuitive evaluation* of any potential action. Considerable research has shown that intuitive evaluations, while fast and efficient, are highly susceptible to affective responses and biased perceptions (Kahneman 2011). This susceptibility extends to risky contexts, in which decision makers often decide based on affective responses (Loewenstein et al. 2001) and, indeed, appear to evaluate risk using their affective responses as a proxy, a strategy known as the *affect heuristic* (Finucane et al. 2000, Slovic et al. 2007).

Why might decision makers' intuitive evaluations overemphasize weaker-links in particular? We reason that, when managing conjunctive risk, decision makers must assess all the barriers to success (i.e., the weaker-link and stronger-link). The relative severity of these barriers will thus become salient to them. Even though the probability of overcoming any given barrier is not intrinsically positive or negative, a weaker-link will be seen in negative terms relative to a stronger-link because it is the more-severe barrier (as in reference dependence, e.g., Kahneman and Tversky 1979). Given decision makers' negative perception of the weaker-link, we argue that they will be susceptible to *negativity bias*—that is, the tendency to view and experience negative entities as more impactful than neutral or positive entities (Rozin and Royzman 2001). If decision makers succumb to negativity bias, then the weaker-link will tend to dominate several aspects of their thinking, including how they direct their attention, their affective response to possible actions they could take, and their perceptions of how impactful those actions might be. The weaker-link will thus take on exaggerated importance in intuitive evaluations of any potential intervention. In sum, and consistent with the tendency to rely on relatively few, key inputs to handle complex decisions (Tversky and Kahneman 1974, Gigerenzer and Goldstein 1996, Kahneman and Frederick 2002, Gigerenzer and Gaissmaier 2011), when facing the complex task of managing conjunctive risk, decision makers may thus primarily rely on the impact of an improvement *on the weaker-link*. We term this proposed decision process the *worst-first heuristic*.

Such a worst-first heuristic is consistent with our empirical results. First, in Studies 1 and 2, we establish the bias toward improving weaker-links. We find that participants are more likely to expend effort to improve weaker-links than stronger-links (Study 1). Moreover, this effect manifests even when improving stronger-links yields greater aggregate chances of success, thereby violating axiomatic rational choice (Study 2).

In Studies 3–5 and Studies S1–S3 in the online supplement (see Table A.1 in the appendix for the table of contents of the online supplement), we present evidence regarding how this bias emerges. This evidence is consistent with negativity bias, which implicates both judgment-based processes and affect-based processes. In line with negative aspects of a situation having more impact on decision makers' intuitive evaluations than neutral or positive aspects, we find clear evidence that participants judge weaker-link improvements to have a larger effect on aggregate chances of success than stronger-link improvements (Study S1), and they cite such judgments as their reason for improving weaker-links rather than stronger-links (Study S2). In line with participants experiencing greater affective reactions to more-negative aspects of a situation and a high level of confidence in their intuitive evaluations (Simmons and Nelson 2006), the effect remained when we gave participants information that should have rendered their intuitive judgments irrelevant. Specifically, participants improved weaker-links more than stronger-links even when we computed and displayed the aggregate consequences of improvements for participants in separate evaluation (Study 3) or told them that the improvements in weaker-links and stronger-links yielded the same aggregate chance of success in joint evaluation (Study S2). Together, these results are consistent with two processes. First, participants might have affective reactions that influence their preferences toward addressing weaker-links, and, second, they might misjudge the actual impact of improvements on their overall chances of success. However, our studies do not show which of these processes is most important in generating the worst-first heuristic or how these processes combine and influence one another. We propose that, consistent with negativity bias, both affective responses and misjudgments of impact play a role (for past research on how affect and cognition play a joint role in decision making, see Slovic et al. 2007, Weber and Johnson 2009).

We also find evidence that it is specifically *managing* conjunctive risk that leads participants to place exaggerated importance on weaker-links—as opposed to merely *assessing* or *choosing between* separate conjunctive risks. We find this overprioritization of weaker-links when participants must decide whether to improve stronger-links versus weaker-links, but not when the same formal decision is presented as a choice of which conjunctive lottery to enter (Study 4). Why would decision makers use a worst-first heuristic particularly when managing conjunctive risk? We propose the following. When the choice is framed as an opportunity to address one of the obstacles within existing conjunctive risk (i.e., managing risk), it would be natural to compare the relative severity of the obstacles, and this comparison may cause the weaker-link to stand out as the biggest problem. Decision makers will perceive the weaker-link in more-negative terms and so apply the worst-first heuristic. When the same choice is instead framed as being between two completely separate conjunctive lotteries, the focal comparison is between two conjunctive risks with separate sets of obstacles rather than between the obstacles within a given conjunctive risk. Therefore, we would expect the relative severity of the obstacles within each lottery to be less salient and decision makers to be less likely to apply the worst-first heuristic. These considerations should cause more prioritization of weaker-links when managing conjunctive risk than when choosing between conjunctive risks, and the results in Study 4 support this prediction. In addition, these results rule out any explanation based on how decision makers might evaluate separate conjunctive lotteries (Nilsson et al. 2009, Jenny et al. 2014) or based on quirks of the probability weighting function (Gonzalez and Wu 1999).

We also find two more results that are consistent with participants primarily relying on the change in the weaker-link as an intuitive proxy for evaluating improvements in conjunctive risk overall. First, participants protect the weaker-link when the decision is about choosing between *reductions* in chances of success, akin to decisions of budgetary cuts during an ongoing project (Study 5). Second, participants are more responsive to the impact of potential improvements in weaker-links than in stronger-links (Study S3). These results are again consistent with the worst-first heuristic, which entails *sensitivity* to changes in weaker-links as opposed to a strict rule to act on weaker-links wherever possible.

We explore whether the worst-first heuristic might apply in managerial settings in which precise probabilities may not be available (in Study 6). To do so, we use a sample of managers who choose whether to improve a weaker-link or stronger-link in a hypothetical managerial scenario in which risk is communicated nonnumerically. We again find a bias toward improving the weaker-link.

Finally, we return to the alternative prediction that decision makers naively judge improvements in the chances of meeting requirements based on their absolute magnitude, which would bias decision makers toward improvements to the stronger-link. Using data from all studies, we compare the behavior of attentive decision makers to that of less-attentive decision makers. Across all participants, the worst-first heuristic is the dominant tendency, but, among the subset of participants with low attentiveness or miscomprehension, we see indications of bias toward improving stronger-links. This suggests that these less-attentive decision makers may fail to appreciate the complexity involved in evaluating changes in conjunctive risk and thus judge improvements in weaker-links or stronger-links based solely on their absolute magnitude. It also suggests that those who use the worst-first heuristic at least appreciate its complexity but still have difficulty navigating it according to a fully rational, calculative process. They instead rely on an intuitive, heuristic evaluation that reflects negativity bias and thus relies primarily on changes in the weaker-link.

On a practical level, the most fundamental implication of our empirical results is the following: when decision makers manage conjunctive risk, they may overprioritize the biggest barrier to success. This implication is potentially far-reaching, as conjunctive risk exists in a vast array of settings—any in which a decision maker requires more than one probabilistic event to achieve success or avoid failure. For a disease to be cured, the prescribed treatment must be effective and the patient must follow through with whatever treatment he or she is given, both of which are uncertain. For a new product idea to develop into a successful business, an entrepreneur faces both technological risk in developing the innovation and market risk in whether consumer demand for that product will materialize. For firms or societies to survive, they must successfully navigate multiple threats. In all of these cases, decision makers must decide how to invest their efforts to manage the conjunctive risk. Should the doctor prescribe a treatment that is less often effective but more likely to be adhered to? Should the entrepreneur invest in technology development or marketing initiatives? Should a decision makers who are concerned with survival allocate resources to address the most severe threat or spread their resources across multiple threats? And, more generally, how do decision makers make trade-offs across these priorities? In all these cases, our work suggests that decision makers might overemphasize whichever obstacle to success is most severe and neglect opportunities to make significant progress on other critical obstacles.

On a theoretical level, this article sheds new light on the importance of differentiating the psychology of whether to take a new risk from the psychology of managing existing risk. In so doing, it contributes to a burgeoning behavioral literature on how decision makers manage, hedge, and improve risks to which they are exposed (Frederick et al. 2018, Markle and Rottenstreich 2018, Lewis and Simmons 2020). For example, Frederick et al. (2018) find that decision makers' implied risk preferences in their decisions about which risks to hedge (managing risks) are negatively correlated with their implied risk preferences in their decisions about when to take risks (choosing which risks to opt into). Markle and Rottenstreich (2018) find that, whereas decision makers are comparatively reluctant to hedge an existing bet themselves (managing risks), they often prefer a prehedged bet to a riskier bet (choosing which risks to opt into). Finally, Lewis and Simmons (2020) find that decision makers have a greater tendency to improve favorable versus unfavorable probabilities (managing risks) than would be implied by their decisions about separate risky prospects (choosing which risks to opt into). This paper makes an analogous distinction in the context of conjunctive risks: decision makers are biased toward improving weaker-links over stronger-links (managing risks), but this is not reflected in their choices between different conjunctive lotteries (choosing which risks to opt into). This distinction also helps us build on past work that has studied how decision makers generally value conjunctive lotteries, but which has not examined the psychology of *managing* conjunctive risk (e.g., Nilsson et al. 2009, Jenny et al. 2014). When decision makers must manage conjunctive risk and determine which obstacles to address, they will be more likely to consider which obstacle is most severe, leading them to see the weaker-link in negative terms and place exaggerated weight on it.[2]

## Research Methodology

Across nine online experiments, including three in the online supplement (total $n$ = 3,653), we accumulate evidence consistent with the use of a worst-first heuristic. For robustness, we implement multiple experimental paradigms across these studies. We find a bias toward prioritizing weaker-links in both separate evaluation (either deciding whether to effortfully attain an improvement in a weaker-link or, separately, deciding whether to effortfully attain an improvement in a stronger-link; Studies 1–3, Study S3) and joint evaluation (deciding between an improvement in a weaker-link and an improvement in a stronger-link; Studies 4–6, Studies S1–S2). In the separate evaluation studies, to ensure adequate statistical power, participants made multiple decisions and we manipulated within subjects whether participants could improve the weaker-link or stronger-link in each decision. In Online Supplement 8, we present evidence for the effect on a between-subjects basis.

The sample size, exclusions, and primary analyses for all studies reported in this paper were preregistered on aspredicted.org. All materials and data are available at https://researchbox.org/155, and the preregistration links are in the appendix (Table A.2). Any deviations from our preregistrations are noted in the study sections and have no effect on significance at alpha levels of 0.001, 0.01, or 0.05. In Studies 1, 2, 5, and S3, we preregistered exclusions based on the time that participants spend on various survey questions, but the bias toward improving weaker-links remains statistically significant at alpha levels of 0.05 for all of these studies if we remove this exclusion. We report ordinary least squares (OLS) regressions with binary dependent variables, but significance does not change when using logit models instead (see Online Supplement 14).

## Study 1: A Greater Tendency to Improve Weaker-Links

Participants either decided whether to invest effort to improve the chance of the less-likely requirement (the weaker-link) or made the same decision for the more-likely requirement (the stronger-link).

### Method

**Participants.** We aimed to recruit 400 U.S. MTurk participants (before exclusions) for $1.00 each. All sample sizes for Studies 1–5 were determined based on what we could sustainably afford with a minimum of 200 participants per experimental condition, and the sample populations in these studies were chosen for convenience. After preregistered exclusions (https://aspredicted.org/n4m77.pdf) for incomplete entries, duplicate entries, or failing comprehension checks,[3] our final sample was $n$ = 252, $M_{age}$ = 39.8, $P_{female}$ = 42% (see Online Supplement 4 for complete details regarding preregistration and the sample). Due to the large number of exclusions for inattentiveness in each of our studies, we analyze the implications of these exclusions after Study 6 in the section on excluded participants, which also yields further theoretical insight.

**Design.** Participants made incentive-compatible decisions about each of eight conjunctive lotteries in which they would win a bonus payment if two independent, probabilistic "requirements" were met (Requirement 1 and Requirement 2). For example, for the conjunctive lottery in Figure 1(a), the participant would win $0.75 if and only if both Requirement 1 was met (which has a 40% chance) and Requirement 2 was met (which has a 20% chance). Participants could decide whether to improve the chance of Requirement 1 being met by investing effort through typing "ab" on a keyboard a given number of times (DellaVigna and Pope 2018). For example, if faced with the lottery in Figure 1(a), a participant could improve the chance of Requirement 1 being met from 40% to 50% by typing "ab" on a keyboard 45 times.

Within subjects, and across these eight decisions, we manipulated whether the improvable requirement had the higher probability (improve-stronger-link condition) or the lower probability (improve-weaker-link condition). The improvable requirement was always labeled as Requirement 1. For example, Figure 1(a) represents the improve-stronger-link condition, since the improvable 40% chance that Requirement 1 is met is the stronger-link (i.e., it is greater than Requirement 2's 20% chance of being met). Figure 1(b), in contrast, represents the improve-weaker-link condition, since the improvable 20% chance that Requirement 1 is met is the weaker-link.

**Stimuli Generation.** We designed the stimuli such that each participant would make eight effort-for-improvement decisions in total. To see whether participants would make inconsistent choices across two *formally equivalent* situations, these eight decisions were split into four pairs, each including one decision in the improve-weaker-link condition and one formally equivalent decision in the improve-stronger-link condition (see Table 1 for an example of the stimuli that one participant saw at different points in the study). The key to understanding this design is that, for the two decisions in each pair, participants would see the same numbers for (a) the probability of the less-likely requirement being met (the weaker-link), (b) the probability of the more-likely requirement being met (the stronger-link), (c) the bonus amount that the participant would receive if both requirements were met, and (d) the amount of work that the participant would have to do to improve their chances of winning that bonus. Each pair of decisions with these same numbers would include one decision about whether to invest the effort to improve the lower-probability requirement (improve-weaker-link condition; see Figure 1(a)) and one about whether to invest the effort to improve the higher-probability requirement (improve-stronger-link condition; see Figure 1(a)). The magnitudes of the improvements in the stronger-link and weaker-link were set such that they would have an identical effect on the aggregate chance of winning the bonus money. Because the decisions within each pair were formally equivalent, any condition-differences in behavior within the same pair would be inconsistent.

In order to (a) ensure that any result holds for a range of possible percentages and bonus amounts, (b) be transparent about how the numbers were chosen, and (c) avoid inadvertently selecting numbers for which left-digit bias might exaggerate the potential improvement in either the weaker-link or stronger-link (Poltrock and Schwartz 1984), we randomly sampled the parameters in the following way. For each pair, we randomly drew the potential bonus prize from $0.50 to $1.00, the initial probability of the stronger-link from integers between 40% and 60%, the initial probability of the weaker-link from integers between 10% and 30%, and the possible improvement in the stronger-link from integers between 10% and 39%. The possible improvement in the weaker-link was then set to be equivalent to the possible improvement in the stronger-link (i.e., to result in the same aggregate probability of winning the bonus) using the following formula (rounded to the nearest percentage point)[4]:

**Figure 1.** Study 1 Screenshots

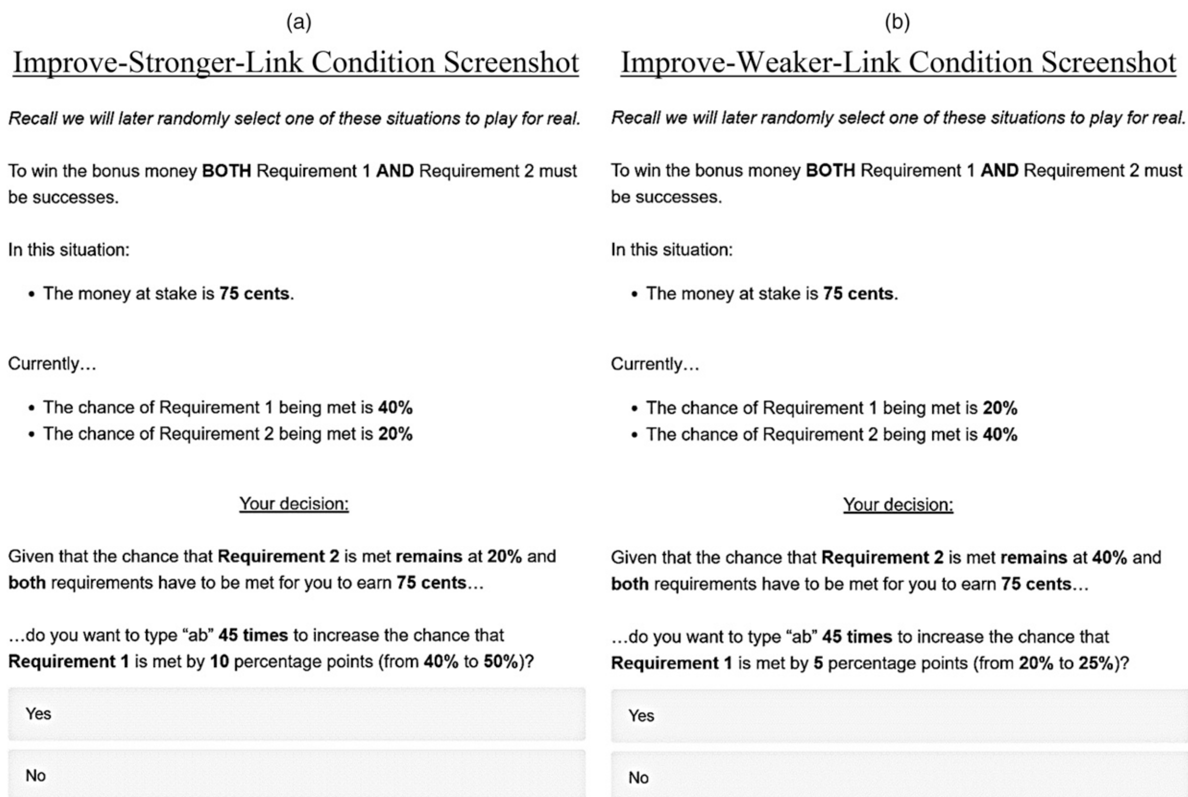


*Notes.* Figure 1 displays a screenshot for each of two examples of decision pages from Study 1. Panel (a) shows a decision in the improve-stronger-link condition, and panel (b) shows the equivalent decision in the improve-weaker-link condition. The initial probability of the weaker-link is 20%, the initial probability of the stronger-link is 40%, and the required amount of typing to attain that improvement is 45 "ab's". The possible improvement in the stronger-link is 10 percentage points (from 40% to 50%; see panel (a)) and the possible improvement in the weaker-link is five percentage points (from 20% to 25%; see panel (b)). In both cases, the aggregate chance of success postimprovement is 10%.

$$= \textit{possible improvement in stronger-link} \times \frac{\textit{initial weaker-link}}{\textit{initial stronger-link}}.$$

Lastly, the amount of work on offer to attain a given improvement was set such that the participant had to type "ab" a number of times equal to 30 × the expected value gain from the potential improvement (in cents, with this product rounded to the nearest whole number). See Table 1 for an example of a full set of stimuli that one participant saw and how the order of decisions was randomized without both decisions in any one pair appearing next to each other.

**Procedure.** At the beginning of the survey, to give participants some intuition for how effortful the "ab" typing task would be, we asked them to practice typing "ab" and allowed them to continue to the next page after they had typed "ab" 50 times. We then explained to them that they would make decisions about lotteries in which they would win a bonus payment if two independent, probabilistic requirements were met. Participants were told that they could decide whether to improve the chance of one of the requirements being met by doing a typing task like the one they had done at the beginning of the survey. To ensure that participants took their decisions seriously, we truthfully told them that we would randomly select one of these eight lotteries to conduct for real. Participants were then given an example decision and asked four comprehension questions about it. These comprehension checks tested whether they understood (a) what would happen if they agreed to the typing task for the decision selected to count, (b) what would happen if they did not agree to do the typing task for this decision, (c) that they needed both requirements to be met to win the bonus, and (d) that the requirements were statistically independent of each other (see Online Supplement 5 for details of comprehension questions).[5] Participants had two opportunities to answer correctly—but we only included participants in the main analysis who answered questions (a), (c), and (d) correctly on their first attempt. We had preregistered to insist on participants passing (b) on their first attempt as well, but a minor typographical error in the instructions made

**Table 1.** Example of Randomly Generated Stimuli for a Participant in Study 1

| Pair | Requirement 1 (Initial probability) | Requirement 2 (Unimprovable probability) | Possible improvement to chance of Requirement 1 (Nonrounded improvement precisely equivalent to that of stronger-link) | Improved probability of Requirement 1 | Final improved aggregate chance of success (not displayed to participants) | Bonus | Required amount of typing "ab" | Condition | Display order |
|---|---|---|---|---|---|---|---|---|---|
| 1 | 58% | 14% | 18 percentage points | 76% | 10.6% | $0.74 | 56 | Improve stronger-link | 1 |
| | 14% | 58% | 4 (4.344 . . .) percentage points | 18% | 10.4% | | | Improve weaker-link | 6 |
| 2 | 57% | 27% | 13 percentage points | 70% | 18.9% | $0.73 | 77 | Improve stronger-link | 3 |
| | 27% | 57% | 6 (6.157 . . .) percentage points | 33% | 18.8% | | | Improve weaker-link | 7 |
| 3 | 52% | 29% | 14 percentage points | 66% | 19.1% | $0.71 | 86 | Improve stronger-link | 5 |
| | 29% | 52% | 8 (7.807 . . .) percentage points | 37% | 19.2% | | | Improve weaker-link | 2 |
| 4 | 57% | 20% | 35 percentage points | 92% | 18.4% | $0.99 | 208 | Improve stronger-link | 8 |
| | 20% | 57% | 12 (12.280 . . .) percentage points | 32% | 18.2% | | | Improve weaker-link | 4 |

*Notes.* All information corresponds to the stimuli generated for a particular participant in Study 1. Each pair includes two decisions that share the same initial stronger-link and weaker-link. Which of the stronger-link or weaker-link can be improved (i.e., which one is Requirement 1) depends on the condition. The bonus and the required amount of typing is the same for the decisions in each pair, and the improvements in each pair have approximately the same impact on the aggregate change of success.

this question unreasonably difficult (see Online Supplement 5). This deviation from our preregistration does not change the significance of the main effect of the improve-weaker-link condition at an alpha level of 0.001.

Participants then made their decisions, which were presented one at a time on the computer screen. To prevent participants from realizing that some of their decisions were economically equivalent, we randomized the order of decisions so that there were always at least two additional decisions splitting each pair of economically equivalent improve-weaker-link versus improve-stronger-link decisions (Table 1 above has a column detailing how the display order was randomized). Figure 1 above shows screenshots of what participants would have seen for one possible pair of decisions.

After participants made their decisions, the survey randomly selected one of the lotteries to conduct for real and informed them of which lottery had been selected. If the participant previously agreed to type "ab" the necessary number of times to improve this lottery, then they were required to do so before they could complete the survey and were then informed whether they won (with the improved probability of meeting Requirement 1). If the participant previously declined to type "ab" the required number of times for the selected lottery, then they immediately found out (i.e., without doing any extra typing) whether they won with the initial probability of Requirement 1.

At the end of the survey, participants entered demographics information. We paid participants as promised after the study had completed.

## Results

Each participant provided eight observations. The dependent variable measured whether, for each decision, the participant decided to type "ab" the required number of times to increase the probability of Requirement 1 (1 = yes, 0 = no). To test whether decision makers would be more likely to invest effort to increase the chance of meeting the lower-probability requirement (the weaker-link) or the chance of meeting the higher-probability requirement (the stronger-link), we used OLS to regress the dependent variable on an improve-weaker-link condition variable. To account for nonindependence of observations from the same individual, we clustered standard errors by participant. (The preregistered regression also includes an interaction with the ratio of the improvements in the stronger-link and weaker-link, but we report this simpler regression for ease of exposition. It has no effect on the statistical significance of the improve-weaker-link condition variable. See Online Supplement 14 for the exact details of all preregistered and reported regressions in this article, including replications with logit models.)

Participants were 10 percentage-points more likely to invest effort to improve the weaker-link (59%) than they were to improve the stronger-link (49%) in formally equivalent decisions ($b = 0.096$, clustered $SE = 0.023$, $p < 0.001$). As mentioned in the introduction, the direction of this difference was opposite of our initial prediction that decision makers would improve stronger-links more-often based on the larger absolute magnitudes of those improvements. Instead, although improvements in stronger-links and weaker-links were formally equivalent, and although the improvements in stronger-links were of greater absolute magnitude, participants were more likely to expend effort to improve weaker-links than stronger-links.

Across our subsequent studies, we developed and tested an explanation for this effect: decision makers use a worst-first heuristic to manage conjunctive risk, whereby they intuitively evaluate improvements according to their impact on the weaker-link. Study 2 tests an implication of this worst-first heuristic: decision makers will improve weaker-links more than stronger-links even when improvements in stronger-links yield superior chances of a payoff.

## Study 2: Violating Rational Choice Theory

This study tested whether, consistent with the worst-first heuristic, decision makers would exhibit a bias toward improving the weaker-link, even when improving the stronger-link would result in a greater aggregate chance of success (rather than the *same* aggregate chance of success as in Study 1). Such a result would violate rational choice theory.

### Method

**Participants.** We aimed to recruit 600 U.S. MTurk participants (before exclusions) for $1.00 each. After preregistered exclusions (https://aspredicted.org/fy4ed.pdf) for incomplete entries, duplicate entries, failing comprehension checks, and rushed entries (defined by taking the minimum amount of time spent on a decision by each participant and excluding the 25% of participants with the lowest minimum times), our final sample was $n = 248$, $M_{age} = 38.6$, $P_{female} = 48\%$ (see Online Supplement 4 for details of how the final sample was determined). We analyze the implications of this large number of exclusions after Study 6 in the section on excluded participants.

**Design.** The design of Study 2 was largely identical to the design of Study 1, in that we manipulated the improve-weaker-link condition versus the improve-stronger-link condition within subjects; however, there was one important change. In this study, the possible improvement in the stronger-links yielded a *greater* increase in the aggregate chance of winning the bonus than the possible improvement in the weaker-links.

**Stimuli Generation.** To ensure that improving stronger-links yielded greater overall chances of success than improving weaker-links, we started by randomly sampling the initial probabilities of the stronger-links, the initial probabilities of the weaker-links, and the bonus prizes in the exact same way as in Study 1. Then, rather than simply generating the possible improvement in the stronger-link as we had done in Study 1, we generated a *preadjustment* improvement in the stronger-link (this time from between 10 percentage points and 30 percentage points), and then calculated the corresponding improvement in the weaker-link to have the same impact on the aggregate probability of winning the bonus as that preadjustment improvement in the stronger-link would have had. However, we then made the critical departure from the stimuli generation procedure of Study 1. To generate a possible improvement in the stronger-link that would be *more* valuable than the possible improvement in the weaker-link, we added either 1, 2, 3, or 4 percentage points to the preadjustment improvement in the stronger-link. For example, if we initially generated a preadjustment improvement in the stronger-link of 30 percentage points (and so set the potential improvement in the weaker-link to have the same impact on the aggregate chance of a payoff as that 30-percentage-point improvement in the stronger-link), we would then make the potential improvement in the stronger-link either 31, 32, 33, or 34 percentage points (instead of 30) to ensure that improving the stronger-link was objectively more valuable than improving the weaker-link.

Finally, to ensure that improving stronger-links really was more worthwhile than improving the weaker-links, we required participants to exert the same amount of effort to attain the more-valuable improvements in stronger-links and the less-valuable improvements in weaker-links (i.e., the same amount of "ab" typing was required to earn the improvement across the yoked pair of possible improvements, as in the previous study; see Online Supplement 10 for details).

**Procedure.** The procedure of Study 2 was identical to that of Study 1, except for the experimental design differences previously described.

### Results

As in Study 1, each participant provided eight observations. The dependent variable measured whether the participant decided to improve their probability of winning the bonus (1 = yes, 0 = no). To test whether participants would be more likely to invest effort for an improvement in the weaker-link than a more-valuable improvement in the stronger-link, we preregistered an OLS regression of the dependent variable on four dummy variables: (1) a dummy variable for the improve-stronger-link condition when one percentage point had been added to the preadjustment improvement, (2) a dummy variable for the improve-stronger-link condition when two percentage points had been added to the preadjustment improvement, (3) a dummy variable for the improve-stronger-link condition when three percentage points had been added to the preadjustment improvement, and (4) a dummy variable for the improve-stronger-link condition when four percentage points had been added to the preadjustment improvement. We clustered standard errors by participant. Participants were less likely to improve stronger-links than weaker-links, even when the improvement in the stronger-link was boosted by one percentage point ($b = -0.054$, clustered $SE = 0.025$, $p = 0.029$), two percentage points ($b = -0.083$, clustered $SE = 0.028$, $p = 0.004$), three 3 percentage points ($b = -0.091$, clustered $SE = 0.027$, $p < 0.001$), and four percentage points ($b = -0.083$, clustered $SE = 0.027$, $p = 0.002$; see Figure 2 for the percentage choosing to improve in each case).

In sum, participants were more likely to expend effort to improve weaker-links than stronger-links, even though improving stronger-links yielded greater aggregate chances of success.

## Study 3: The Weaker-Link Bias Remains When Decision Makers Know Aggregate Probabilities

This study was designed to test whether the bias toward improving weaker-links would remain when decision makers knew the impact of a potential improvement on the aggregate chance of success. Specifically, for each of their decisions in this study, participants had the potential aggregate improvements in their chance of success computed and shown to them. According to the worst-first heuristic, decision makers intuitively evaluate actions according to their effect on the weaker-link, in part because of negativity bias. Since negativity bias would cause strong intuitions, these intuitions might be sufficient to influence how decision makers react to a known change in their aggregate chance of success (Simmons and Nelson 2006). Moreover, since negativity bias involves affect as well as perceptions (Rozin and Royzman 2001), we would expect decision makers to feel a greater urge to improve weaker-links than stronger-links, even when they knew the resulting aggregate chance of success would be the same in either case. In sum, our account of the worst-first heuristic predicts a bias toward improving weaker-links, even when decision makers have information about aggregate chances of success.

### Method

**Participants.** We aimed to recruit (before exclusions) 500 U.K. and U.S. participants from Prolific for $0.90 each. After preregistered exclusions (https://aspredicted.org/gb9iz.pdf) for incomplete entries, duplicate entries, failing comprehension checks, rushed entries (defined by taking the minimum amount of time spent on a decision by each participant and excluding the 25% of participants with the lowest minimum times), our final sample was $n = 240$, $M_{age} = 33.2$, $P_{female} = 55\%$, $P_{U.K.} = 48\%$ (see Online Supplement 4 for details of how the final sample was determined). We analyze the implications of this large number of exclusions after Study 6 in the section on excluded participants.

**Design.** The design of Study 3 was largely identical to the design of Study 1, employing a two-cell improve-weaker-link versus improve-stronger-link within-subjects design. There was, however, one important change: we told participants what the improvement in the stronger-link or weaker-link would mean for their aggregate chance of winning the bonus for each of their decisions.

**Stimuli Generation.** Stimuli were generated in the exact same way as in Study 1, except that the potential bonus prize was always $0.50.

**Procedure.** The procedure of Study 3 was identical to that of Study 1, with a few slight changes. Primarily, we reworded the instructions and the pages on which participants made their decisions to account for the fact that they would learn the improvement in their aggregate chance of success. To ensure that participants were aware that this information was available and that they could use it to make their decisions, we added an extra comprehension question. They had to answer this question by using the information about the aggregate chance of success in an example decision (see Online Supplement 5 for full details). As preregistered, we only included participants who answered all five comprehension questions (including this extra one) on their first attempt.

Figure 3 shows screenshots of what participants would have seen for one possible pair of decisions. After describing the implications of the improvement for the stronger-link and the weaker-link, we told them, "This would cause the overall chance of winning the lottery to increase by X percentage points (from Y% to Z%)," where X, Y, and Z, were calculated to reflect the true improvement in the aggregate chance of winning the bonus to the nearest percentage point.

As an extra precaution against inattentive participants, we made one additional change. At the beginning of the survey, participants were told that they would make some decisions about improving probabilities. On the next page, they answered an attention-check question in which they had to confirm that the

**Figure 2.** Study 2 Participants Were More Likely to Effortfully Acquire Improvements in Weaker-Links Than More-Valuable Improvements in Stronger-Links

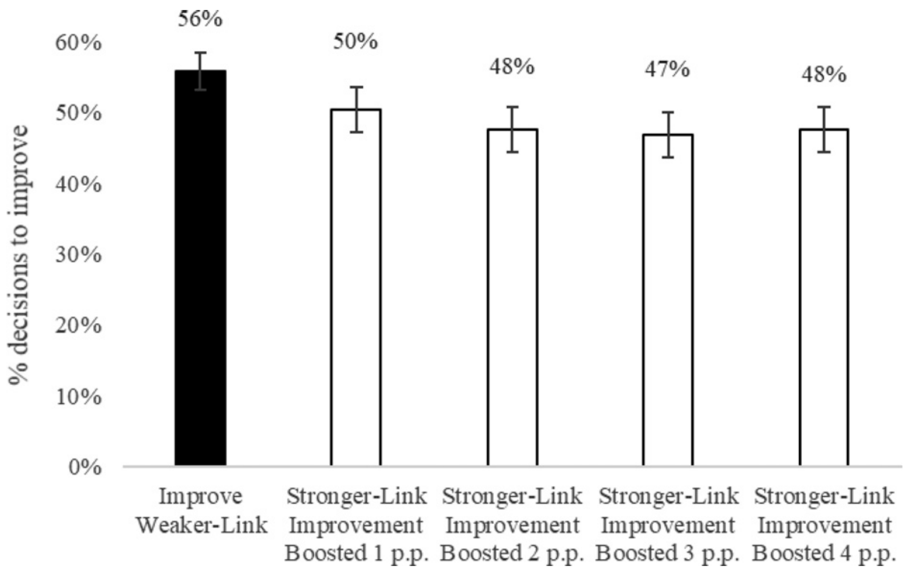


*Notes.* This chart shows the percentage of decisions for which participants agreed to type "ab" the required number of times in order to increase their probability of winning a bonus as a function of whether they could improve the weaker-link (black bar) or the stronger-link by an amount that was one, two, three, or four percentage points (p.p.) higher than the improvement in the stronger-link that would be equivalent to the improvement in the weaker-link. Error bars show ± 1 standard error (clustered standard error in the case of the black bar).

survey was about "improving probabilities" from a list of five options (see Online Supplement 5 for details). Participants who answered incorrectly were prevented from continuing with the survey.

### Results

As in Study 1, each participant provided eight observations. The dependent variable measured whether the participants decided to improve their probability of winning the bonus (1 = yes, 0 = no). We preregistered to use the exact same regression that we reported in Study 1. Even when participants could see the impact that improving the weaker-link or the stronger-link would have on their aggregate chance of winning a bonus, they still improved weaker-links more than they improved stronger-links ($b = 0.048$, clustered $SE = 0.017$, $p = 0.006$). This result is consistent with the proposed role of negativity bias, which would lead participants to have both strong intuitions and affective responses that would each influence how they evaluate a known improvement in the aggregate chance of success. The affective responses associated with the worst-first heuristic are underlined by a similar result in Study S2 (see Online Supplement), in which participants must choose between improving a weaker-link, a stronger-link, or accepting a bonus in return for leaving which improvement they get to chance. Even when told that their aggregate chance of success would be the same no matter which improvement they chose, participants still chose to improve weaker-links more than stronger-links.

## Study 4: Managing Conjunctive Risk vs. Assessing Whether to Enter It

Studies 1–3 all addressed the case of managing existing conjunctive risk. This setting was designed as a proxy for the managerial situation in which an ongoing project requires the achievement of multiple subgoals for ultimate success and the manager must decide whether and how to invest resources in achieving each subgoal. However, not all managerial decisions about conjunctive risk are of this nature. Sometimes, managers must simply decide which projects to undertake and assess the conjunctive risks inherent to each potential project on a take-it-or-leave-it basis. Thus, rather than managing existing conjunctive risk, these decisions are about which conjunctive risks to take in the first place. This distinction also raises an important question about our results thus far: Is excessively prioritizing the weaker-link merely a manifestation of how decision makers think about conjunctive risk in general, or does managing conjunctive risk trigger a distinct psychological process?

In Study 4, we aim to answer this question in an incentive-compatible conjunctive lottery context. To do so, we convert choices between *improving* the stronger-link or the weaker-link (i.e., managing risk) into formally equivalent choices between *entering* the lotteries that would have resulted from improving the

**Figure 3.** Study 3 Screenshots

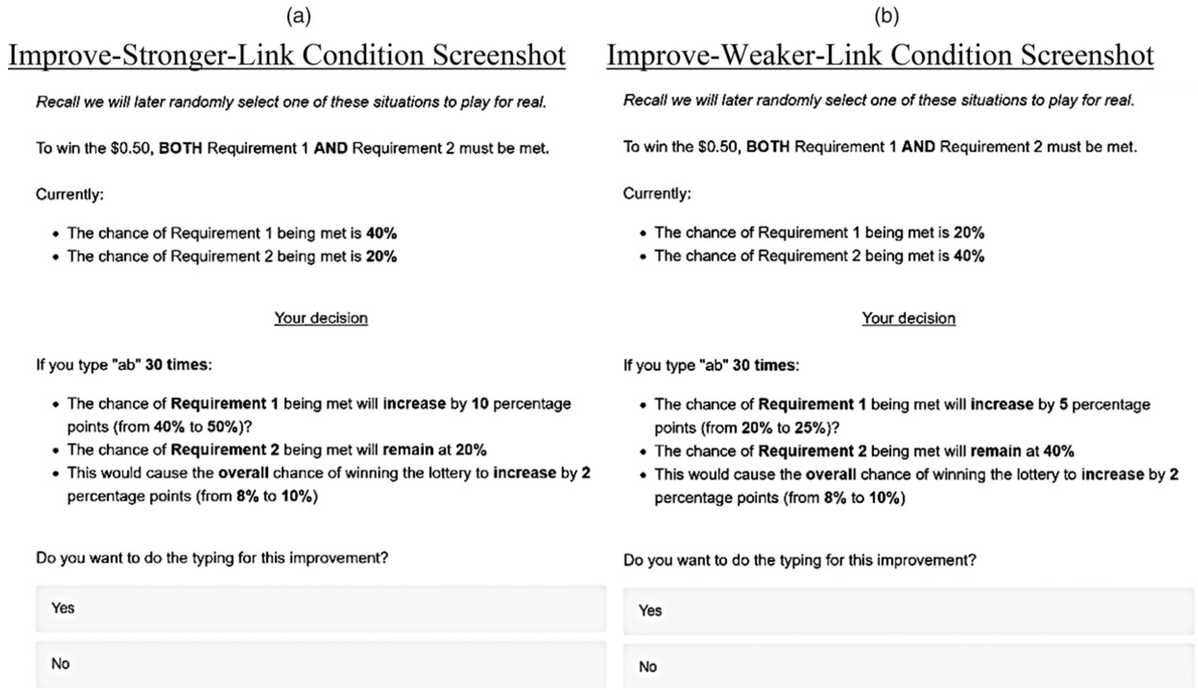


*Notes.* Figure 3 displays a screenshot for each of two examples of decision pages from Study 3. Panel (a) shows a decision in the improve-stronger-link condition, and panel (b) shows the equivalent decision in the improve-weaker-link condition. The initial probability of the weaker-link is 20%, the initial probability of the stronger-link is 40%, and the required amount of typing to attain that improvement is 30 "ab's". The possible improvement in the stronger-link is 10 percentage points (from 40% to 50%; see panel (a)), and the possible improvement in the weaker-link is five percentage points (from 20% to 25%; see panel (b)). In both cases, the aggregate chance of success increases from 8% to 10%, as described in the final bullet point in each panel.

stronger-link or the weaker-link (see the example in Figure 4). According to the logic of a worst-first heuristic, and its requirement that the weaker-link be seen in negative terms, we would expect decision makers to exhibit a stronger bias toward improving weaker-links when the question of which requirement represented the biggest barrier to success was more salient. This question will be more salient when decision makers are choosing whether to *improve* the stronger-link or the weaker-link and are therefore considering which requirement constitutes the worst threat to success. The question will be less salient when decision makers are simply choosing which of two conjunctive lotteries to *enter*, which elicits consideration of each lottery more holistically—this latter choice is, after all, between lotteries rather than whether to improve the weaker-link or stronger-link. As such, decision makers would be more likely to exhibit negativity bias with respect to the weaker-link and, therefore, be more likely to apply the worst-first heuristic in improvement decisions. In contrast, if the weaker-link bias is entirely due to alternative accounts, such as decision makers placing more weight on the lower-probability requirement when evaluating conjunctive lotteries in general (Juslin et al. 2009, Nilsson et al. 2009), some quirk of probability weighting (Gonzalez and Wu 1999), or a general preference for balanced probabilities in line with naive diversification (Benartzi and Thaler 2001) or inequity aversion (Yaari and Bar-Hillel 1984, Hoffman and Spitzer 1985), we should expect a similar weaker-link bias regardless of how this decision is framed.

## Method

**Participants.** We aimed to recruit (before exclusions) 1,000 combined U.K. and U.S. participants from Prolific for $1.00 each. After preregistered exclusions (https://aspredicted.org/af7mr.pdf) for incomplete entries, duplicate entries, or failing comprehension checks, our final sample was $n$ = 840, $M_{age}$ = 33.4, $P_{female}$ = 52%, $P_{U.K.}$ = 51% (see Online Supplement 4 for details of how the final sample was determined). We analyze the implications of this large number of exclusions after Study 6 in the section on excluded participants.

**Design.** Participants were randomly assigned to either an improvement-choice condition or a lottery-choice condition in a between-subjects design. In the improvement-choice condition, participants made a choice about whether to improve the stronger-link or the weaker-link in a conjunctive lottery for a $0.50 bonus payment. In the lottery-choice condition, participants made a mathematically identical choice, but framed as being between two conjunctive lotteries.

**Stimuli Generation.** We generated the parameters in the improvement-choice and lottery-choice conditions so that the decisions in both conditions were ultimately between the same conjunctive lotteries. Accordingly, participants in the lottery-choice condition chose between a lottery that would result from improving the stronger-link and the lottery that would result from improving the weaker-link in the corresponding improvement-choice condition. For example, a participant in the improvement-choice condition might choose between improving the weaker-link from 20% to 28% or improving the stronger-link from 40% to 56%. A participant in the lottery-choice condition might choose between a lottery with a 28%-chance weaker-link and a 40%-chance stronger-link (i.e., the lottery attained from improving the weaker-link to 28%) and a lottery with a 20%-chance weaker-link and a 56%-chance stronger-link (i.e., the lottery attained from improving the stronger-link to 56%).

As in previous studies, we randomly sampled the relevant parameters from predetermined ranges. To generate the improvement choice parameters, we randomly drew the initial probability of the stronger-link from integers between 45% and 80%, the initial probability of the weaker-link from integers between 5% and 40%, and the possible improvement in the stronger-link from integers between 16% and 20%. The possible improvement in the weaker-link was set to be equivalent to the possible improvement in the stronger-link (i.e., to have approximately the same impact on the overall probability of winning the bonus). To ensure that the two possible improvements would be equivalent, we set the improvement in the weaker-link equal to the result of the following formula (rounded to the nearest percentage point):

**Figure 4.** Study 4 Screenshots

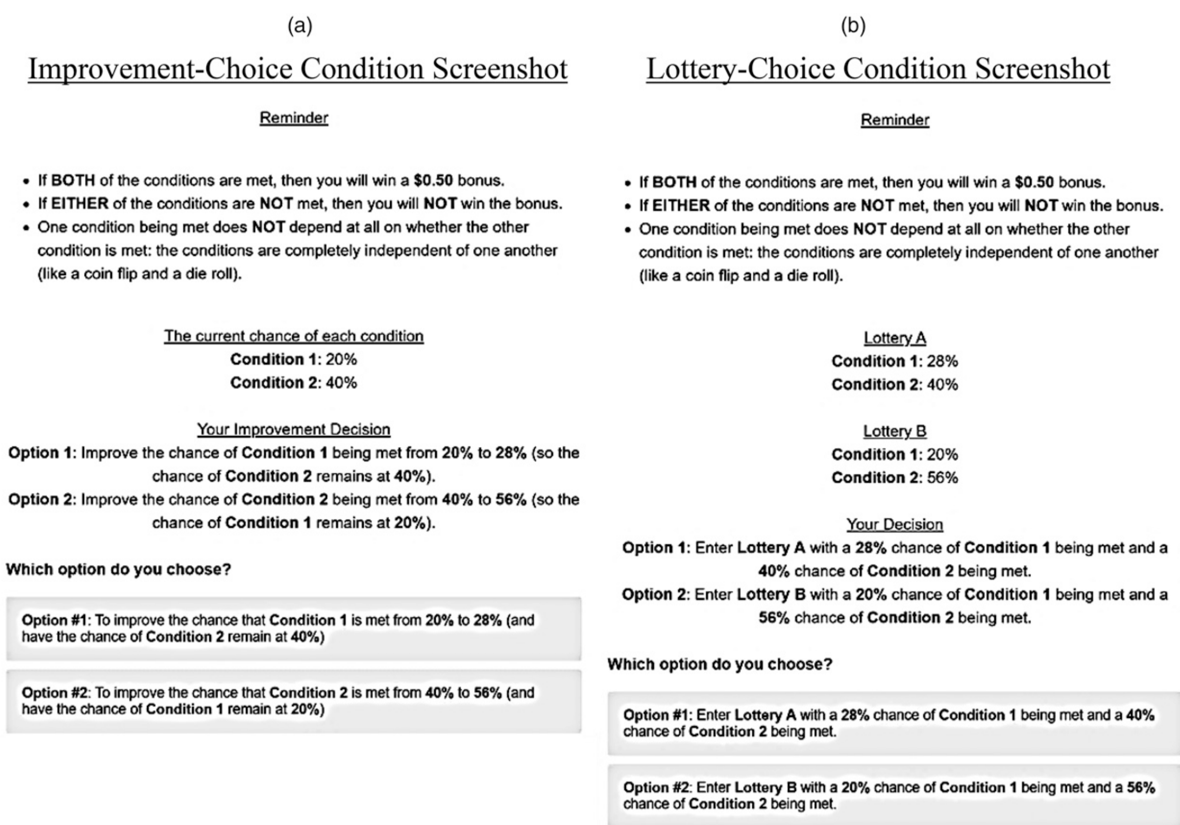


*Notes.* Figure 4 displays screenshots of an example decision from the improvement-choice condition (panel (a)) and the corresponding decision from the lottery-choice condition (panel (b)), which is a formally equivalent decision. The possible improvement in the weaker-link in panel (a) is eight percentage points (from 20% to 28%), and the possible improvement in the stronger-link in panel (a) is 16 percentage points (from 40% to 56%). Therefore, in panel (b), the choice is between the lottery arrived at in panel (a) by improving the weaker-link (28% weaker-link and 40% stronger-link) or the lottery arrived at in panel (a) by improving the stronger-link (20% weaker-link and 56% stronger-link). In these screenshots, the improvement in the weaker-link is presented first, but the order of the options was counterbalanced in the actual study.

$$= \textit{possible improvement in stronger-link} \times \frac{\textit{initial weaker-link}}{\textit{initial stronger-link}}.$$

To generate the lottery-choice condition parameters, we started by generating all parameters in the exact same way as in the improvement-choice condition, and then we used these parameters to formulate the mathematically equivalent lottery choice (as described above). See Figure 4 for an example question, which is formally equivalent across conditions.

**Procedure.** At the beginning of the survey, as a precaution against inattentive participants, we asked participants a similar attention-check question as in Study 3 (see Online Supplement 5 for details). Participants who answered incorrectly were prevented from continuing with the survey. We then explained to participants that they would make decisions about lotteries in which they would win \$0.50 if two independent, probabilistic requirements were met. We asked them two comprehension questions to ensure that they understood these instructions (which we preregistered to use for exclusion purposes). We then explained the nature of their decision according to their assigned condition (improvement choice vs. lottery choice), gave them one decision as an example, and asked them two more comprehension questions about it (which we used for exposition purposes only; see Online Supplement 5 for details).

Participants then made their decision. Figure 4 above shows screenshots of what participants would have seen in the improvement-choice condition decision (panel (a)) and in an equivalent lottery-choice condition decision (panel (b)).

After they made their decision, the survey software conducted the lottery in keeping with their decision. Participants immediately found out whether they had won, and we paid participants bonuses, as promised, upon completion. Participants provided demographic information in the last question of the survey.

## Results

We used a dependent variable that measured whether the participant's choice maximized the weaker-link. In the improvement-choice condition, this dependent variable tracked whether the participant decided to improve the weaker-link or the stronger-link (1 = improve weaker-link, 0 = improve stronger-link), and, in the lottery-choice condition, it tracked whether the participant chose the lottery that was equivalent to improving the weaker-link or to improving the stronger-link (1 = equivalent to improving weaker-link, 0 = equivalent to improving stronger-link). We hypothesized that, because participants would be more likely to consider which requirement was the worst threat to success in the improvement-choice condition, they would be more likely to use a worst-first heuristic in it than in the lottery-choice condition. To test this hypothesis, we preregistered to use OLS to regress the dependent variable on a dummy variable for the improvement condition and a control variable for the improvement ratio between the weaker-link and stronger-link (the improvement in the stronger-link divided by the improvement in the weaker-link). Consistent with the hypothesis, participants were over 14 percentage points more likely to prioritize the weaker-link in the improvement-choice condition ($b = 0.147$, $SE = 0.034$, $p < 0.001$). In the improvement-choice condition, 64% of participants chose to improve the weaker-link (which is significantly more than 50% based on a binomial test, $p < 0.001$). In the lottery-choice condition, however, the bias toward the weaker-link was eliminated. Just 50% of participants chose the lottery that maximized the weaker-link (not significantly different from 50% based on a binomial test, $p = 0.922$). These results (portrayed in Figure 5) are consistent with the notion that decision makers employ a worst-first heuristic, particularly when managing conjunctive risk. They thus illustrate the importance of distinguishing the psychology of managing conjunctive risk from that of deciding which conjunctive risks to take in the first place.

When the ratio between the improvements in the stronger-link and weaker-link is greater (e.g., 15% and 1% as opposed to 9% and 8%, for illustration), the tendency to improve the weaker-link over the stronger-link reduces in magnitude ($b = -0.047$, $SE = 0.011$, $p < 0.001$; see Online Supplement 11 for a discussion of when this relationship might eliminate the bias toward the weaker-link).

**Figure 5.** Study 4 Participants Prioritized the Weaker-Link in the Improvement-Choice Condition but Not in the Lottery-Choice Condition

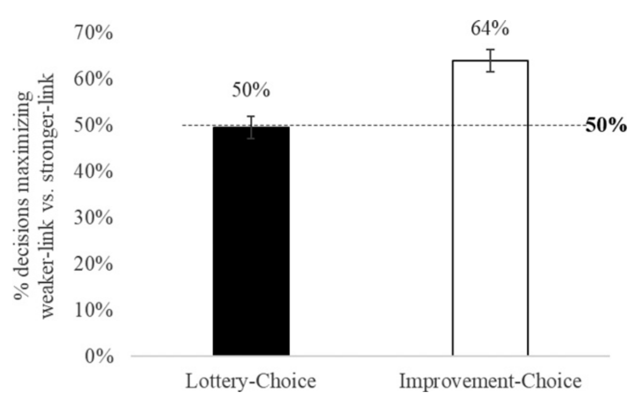


*Notes.* This chart shows the percentage of decisions maximizing the weaker-link rather than the stronger-link, as a function of lottery-choice vs. improvement-choice condition. Error bars show ± 1 standard error.

## Study 5: Choosing Between Reductions in Chances of Success

In organizations, managers often face budget restrictions and must decide where to cut back. When managing conjunctive risk, this might mean accepting a reduction in the probability of meeting a requirement that is necessary for success. If decision makers use a worst-first heuristic, then they will evaluate a reduction in chances of success primarily according to its impact on the weaker-link, and so they will prefer to accept a large reduction in the stronger-link than a small reduction in the weaker-link that is ultimately equally damaging. We tested this prediction by entering participants into a conjunctive lottery and having them make a direct choice between a reduction in the weaker-link and a reduction in the stronger-link.

### Method

**Participants.** We aimed to recruit (before exclusions) 300 combined U.K. and U.S. participants from Prolific for $0.55 each. After preregistered exclusions (https://aspredicted.org/uy6zq.pdf) for incomplete entries, duplicate entries, or failing comprehension checks, our final sample was $n$ = 229, $M_{age}$ = 34.9, $P_{female}$ = 60%, $P_{U.K.}$ = 96% (see Online Supplement 4 for details of how the final sample was determined). We analyze the implications of this large number of exclusions after Study 6 in the section on excluded participants.

**Design.** All participants made a choice about whether to accept a reduction in the stronger-link or the weaker-link (order of options counterbalanced) in a conjunctive lottery for a $0.50 bonus payment.

**Stimuli Generation.** As in previous studies, we randomly sampled the relevant parameters from predetermined ranges. To generate the improvement-choice parameters, we randomly drew the initial probability of the stronger-link from integers between 60% and 80%, the initial probability of the weaker-link from integers between 10% and 30%, and the possible reduction in the stronger-link from integers between 10% and 15%. The possible reduction in the weaker-link was set to be equivalent to the possible reduction in the stronger-link (i.e., to have approximately the same impact on the overall probability of winning the bonus). To ensure that the two possible reductions would be equivalent, we set the reduction in the weaker-link equal to the result of the following formula (rounded to the nearest percentage point):

$$= \textit{possible reduction in the stronger-link} \times \frac{\textit{initial weaker-link}}{\textit{initial stronger-link}}.$$

See Figure 6 for an example question.

**Procedure.** At the beginning of the survey, as a precaution against inattentive participants, we asked participants a similar attention-check question as in Study 3 (see Online Supplement 5 for details). Participants who answered incorrectly were prevented from continuing with the survey. We then explained to participants that they would make decisions about lotteries in which they would win $0.50 if two independent, probabilistic requirements were met. We asked them two comprehension questions to ensure they understood these instructions. We then described one decision as an example and asked them two more comprehension questions about it. We used all four of these comprehension questions as a basis for excluding participants from the analyses (see Online Supplement 5 for details).

Participants then made their decision. Figure 6 shows screenshots of what participants would have seen. After they made their decision, the survey software conducted the lottery in keeping with their decision. Participants immediately found out whether they had won, and we paid participants bonuses, as promised, upon completion. Participants provided demographic information in the last question of the survey.

### Results

We preregistered to test for a bias toward protecting the weaker-link (i.e., reducing the stronger-link) by comparing the proportion of decisions to reduce the weaker-link (vs. the stronger-link) to 50% using a binomial test. Just 41% reduced the weaker-link compared with 59% reducing the stronger-link ($p$ = 0.005).

**Figure 6.** Study 5 Screenshot

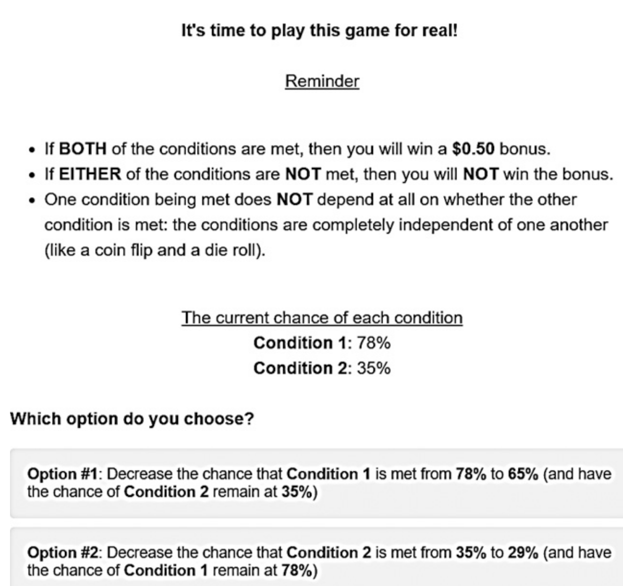


*Notes.* Figure 6 displays a screenshot of an example decision. In this screenshot, the reduction in the stronger-link is presented first, but the order of the options was counterbalanced in the actual study.

This result suggests that decision makers use a worst-first heuristic when deciding which reductions in chances of success to accept; they evaluate such reductions according to their impact on the perceived biggest threat to success.

## Study 6: Managerial Sample and Setting with Nonnumeric Risk Conveyance

To explore the generality and robustness of the bias toward improving weaker-links, this study tests the worst-first heuristic with a sample of managers, in a managerially relevant scenario, wherein risk is communicated nonnumerically.

### Method

**Participants.** We aimed to recruit 200 participants (before exclusions) from ROI Rocket for \$6.50 each in exchange for completing a five-minute survey. Participants were screened to be fully employed in a managerial role of any kind for generalizability to the managerial domain. Because of the relatively large cost of recruiting these participants, we conducted our planned analyses after 110 responses, and all these analyses were extremely statistically significant ($p < 0.0001$). Due to the low risk of false positive associated with such low $p$-values, we stopped collecting data before we met our target sample size. After preregistered exclusions (https://aspredicted.org/ym6a9.pdf) for duplicate attempts, failing comprehension checks, estimating the weaker-link to be higher than the stronger-link, and rushed responses (defined by measuring the total amount of time spent on estimating the stronger-link and weaker-link and then excluding the 10% of participants who spent the least amount of time), our final sample was $n = 67$, $M_{age} = 58.1$, $P_{female} = 37\%$ (see Online Supplement 4 for details of how the final sample was determined). We analyze the implications of the exclusions in the section on excluded participants immediately after the Results section for this study.

**Design.** Study 6 employed a realistic managerial scenario in which a company was trying to develop a new module for a self-driving car before a deadline. To be successful, the company needed two components both to be completed: sensors and navigation algorithms. The chance of one of these components being completed on time (the weaker-link) was lower, and the chance of the other component being completed on time (the stronger-link) was higher. We counterbalanced which component was the stronger-link and which was the weaker-link between subjects. To indicate which component had a lower or higher probability of being completed in a realistic way, we described these probabilities verbally to participants.

**Procedure.** At the beginning of the survey, we introduced a scenario about self-driving cars. To motivate participants to read carefully, we told them that we would ask a question about the scenario on the next page. They had to answer this question correctly to avoid having to attempt it again after rereading the scenario for at least another 30 seconds. The scenario read as follows:

> The Self-Driving Car is a vehicle that does not require any human intervention to move from one destination to another. It promises to revolutionize transportation. It also promises to become a huge market opportunity, worth hundreds of billions of dollars in the coming years.
>
> Your company is in charge of producing a critical control module for the self-driving car. Your production contract promises a **\$100m bonus** if you can deliver this module within the next two months.
>
> There are two critical technologies that still need to be finalized: **sensors** (hardware) and **navigation algorithms** (software). They are the last two things that need to be solved before you can launch.
>
> These two technologies are being developed by two independent teams. The chances of succeeding with one have nothing to do with succeeding with the other. Both technologies need to be developed successfully in order for you to launch.
>
> Recently, you met with each team to discuss its development outlook.

On the next page, participants answered a comprehension check in which they confirmed that, for success, "Two different technologies (sensors and navigation algorithms) need to be developed successfully," from four response options (see Online Supplement 5). If participants answered incorrectly, then they had to reread the scenario for at least 30 seconds and attempt the question again. If they answered incorrectly a second time, then they had to reread the scenario for at least another 30 seconds, but, this time, the correct answer to the comprehension check was revealed on the same page. They could then continue to the next page of the survey. However, as preregistered, we excluded from the analyses any participant who did not answer correctly on the first attempt.

On this page, we described meetings first with the team responsible for sensors and then with the team responsible for navigation algorithms, and, through this description, participants learned the likelihood of finalizing each technology in time. In the description of the technology with the lower chance of being completed (i.e., the weaker-link), participants read the following:

> In your conversation with the team, it was clear that they were facing some challenging obstacles. Meeting

> the deadline would be hard for them but is still possible. Based on what you know, your belief is that their chance of developing their technology in time is **pretty unlikely**. You think it **has some chance**.

In the description of the technology with the higher chance of being completed (i.e., the stronger-link), participants read the following:

> In your conversation with the team, it was clear that some hurdles remained. The timeframe would be a challenge but could be doable. Based on everything you know, your belief is that their chance of developing their technology in time to meet the deadline is **somewhat likely**. You think it is **a bit like flipping a coin**.

On the same page, using sliders going from 0 to 100, participants reported their own perception of the probability of the stronger-link and the probability of the weaker-link.

After doing so, participants moved onto the page on which they made their improvement decisions (displayed in Figure 7). For robustness, we randomized whether we piped participants' estimated probabilities back to them as they made their decision (estimated-probabilities condition, Figure 7(a)) or whether we instead repeated the linguistic description of the probabilities (linguistic-description condition, Figure 7(b)).

As in the Study 6 improvement-choice condition, participants had to choose whether to improve the chance of the stronger-link or the weaker-link. The possible improvement in the stronger-link was set so as to go half-way toward 100% from the participant's estimate of the initial stronger-link probability, using the following formula (and then rounded to the nearest percentage point):

$$= \frac{100 - \textit{estimated stronger-link}}{2}.$$

The potential improvement in the weaker-link was calculated to generate an approximately equivalent improvement in the aggregate chance of success, using the following formula (and then rounded to the nearest percentage point):

$$= \textit{nonrounded possible improvement in the weaker link} \times \frac{\textit{estimated weaker-link}}{\textit{esitmated stronger-link}}.$$

For example, in Figure 7, the weaker-link was estimated to be 20%, and the stronger-link was estimated to be 40%. The possible improvement in the stronger-link is 30 percentage points ($= \frac{100-40}{2}$), so the participant would have the option to improve the stronger-link from 40% to 70%. The equivalent improvement in the weaker-link is 15 percentage points ($= 30\% \times \frac{20\%}{40\%}$), so the participant would have the option to improve the weaker-link from 20% to 35%.

To encourage participants to think carefully about their decision and to gain some insight into their thought process, we asked them to explain it on that same page in an open-ended text box. They then indicated their demographics and completed the survey.

### Results

Recall that the two potential improvements would have equivalent effects on the aggregate chance of success. We preregistered to test for a weaker-link bias by comparing the proportion of decisions to improve the weaker-link (vs. the stronger-link) to 50% using binomial tests. In the estimated-probabilities condition, 83% of participants improved the weaker-link, while in the linguistic-description condition, 90% of participants improved the weaker-link. Overall, in both conditions combined, 87% of participants improved the weaker-link (all three proportions were different from 50% at $p < 0.001$). This result suggests that managers may use a worst-first heuristic in realistic managerial situations, even when risks are expressed nonnumerically.

## Do Excluded Participants Exhibit the Opposite Bias?

The dominant tendency among participants who met our preregistered inclusion criteria was to improve weaker-links, consistent with the use of a worst-first heuristic. However, based on previous literature (Slovic 1972, Burson et al. 2009, Kahneman 2011, Shue and Townsend 2021), we would expect some participants to be naive to the fact that a one-percentage-point improvement in the weaker-link has a bigger impact on the aggregate probability of success than a one-percentage-point improvement in the stronger-link, and mistakenly believe that it was always best to choose the greatest increase in the chance of meeting any individual requirement. Such an absolute-magnitude fallacy would bias participants to opt for stronger-link improvements over equally valuable weaker-link improvements. Given that this fallacy would represent a relatively crude psychology that would not require participants even to understand the complex, conjunctive nature of the problem, it might be more common in inattentive participants. Thus, among inattentive participants, the weaker-link bias that we have documented might disappear or reverse.

We tested this possibility in exploratory analyses on the behavior of the participants who failed to meet the inclusion criteria due to inattentiveness. Figure 8 compares the relative tendency to improve weaker-links versus stronger-links in each of our studies among participants who passed versus failed the preregistered exclusion criteria. In all our studies, apart from

**Figure 7.** Study 6 Screenshots

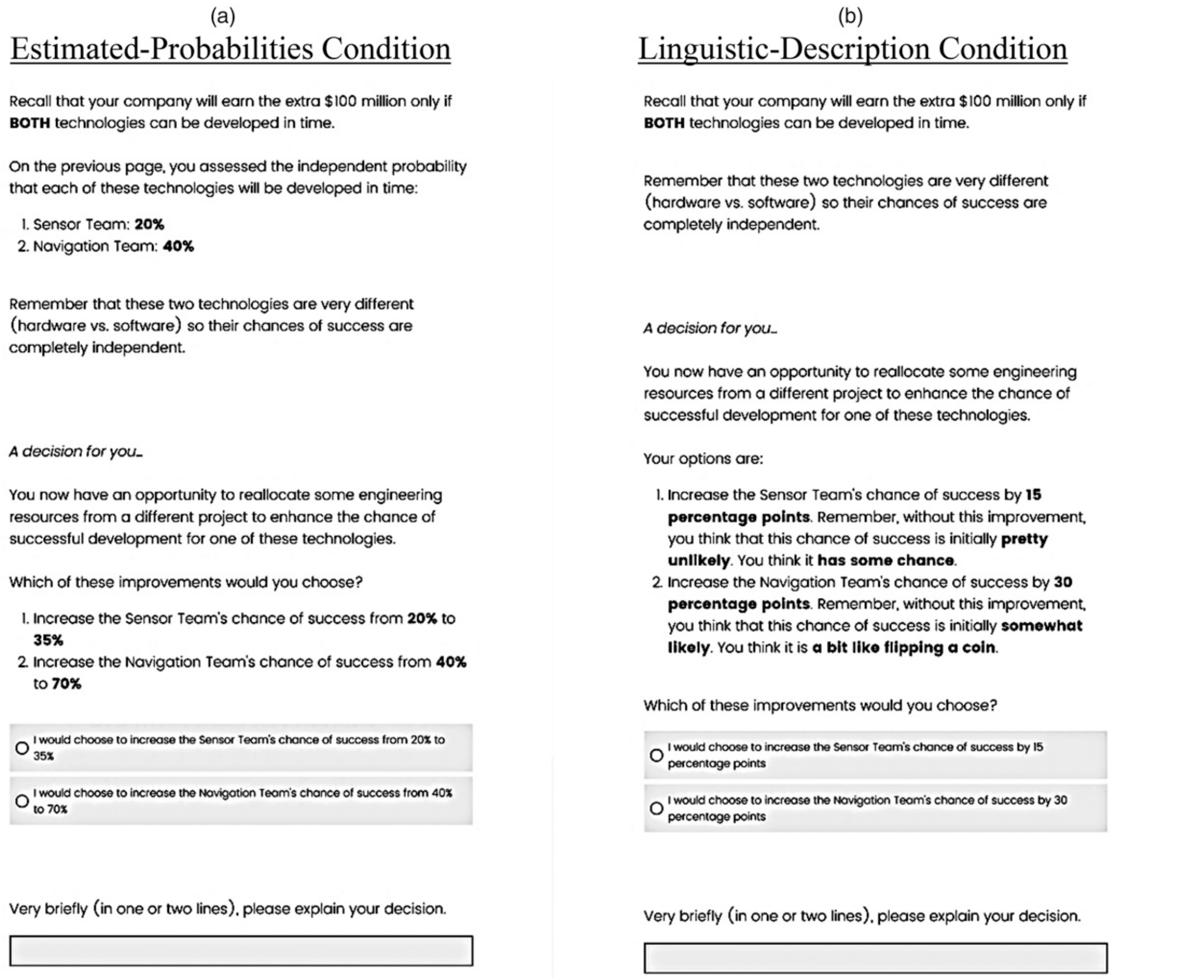


*Notes.* Figure 7 displays screenshots of examples of decisions from the estimated-probabilities condition (panel (a)) and the linguistic-description condition (panel (b)), assuming that the weaker-link was estimated at 20% and the stronger-link at 40%. These screenshots show the improvement in the weaker-link presented first, but the order of the options was counterbalanced.

Study 5, we find that the bias toward improving the weaker-link is greater for participants who passed the preregistered comprehension and attention criteria ($p < 0.012$; see Online Supplement 14 for relevant regressions for Studies 1–4 and S2–S3 and Online Supplement 11 for the relevant $\chi^2$ tests in Studies 5–6, and a $t$-test in Study S1). This result is not particularly remarkable in and of itself, as adding random decision making to a sample will naturally reduce any effect. However, more interestingly, the direction of the bias often flips among excluded participants (resulting in a slight tendency toward improving stronger-links over weaker-links), significantly so in Study 2 (in which this tendency is rational) and Study S3 ($p < 0.001$ in these studies; see Online Supplement 14 for regressions).

It thus appears that many less-attentive participants do judge improvements in stronger-links and weaker-links according to their absolute magnitude and thus exhibit a bias toward the stronger-link. This finding suggests that, indeed, an absolute-magnitude fallacy is a cruder psychology in which decision makers do not appreciate the complex, conjunctive nature of the problem. In contrast, the worst-first heuristic is used most reliably by participants who are paying attention, comprehending the problem, and not rushing through their decision-making process. Still, while these participants are more likely to recognize the complexity of the problem, they still seem to rely on intuitive evaluations rather than rationally calculating and assessing the impact of any improvement in conjunctive risk. Thus, there is a distinction between simple misconceptions on the one hand and, on the other, intuitive decision rules that require greater engagement with the problem yet can likewise cause bias (e.g., Slovic et al. 1965).

Despite this heterogeneity, when the data are pooled by removing the exclusions for inattentiveness, a bias toward improving weaker-links clearly emerges as the dominant tendency. For all studies, we can detect a significant bias toward improving the weaker-link when not applying the time-spent exclusions (see Online Supplement 6 for a deeper analysis of the implications of the time-spent exclusions). For all studies except Studies 2–3, we can detect a significant bias toward improving the weaker-link when not applying any of the inattentiveness and miscomprehension exclusions (at an alpha level of 5%; see Figure 9). Online Supplement 7 has more detail about the prevalence of the bias across participants.

## General Discussion

We studied how decision makers think about managing conjunctive risk—wherein success depends on the occurrence of multiple uncertain events. We found that they more often opted for improvements in the chances of less-likely requirements (weaker-links) than improvements in chances of more-likely requirements (stronger-links) that yielded the same overall chance of success.

We hypothesized that decision makers apply a worst-first heuristic. Given the complexity of precisely calculating all potential aggregate success rates when managing conjunctive risk, decision makers instead form an intuitive evaluation of any potential improvement. Since decision makers identify the weaker-link as their biggest problem, or in terms of negativity bias, perceive it most negatively (Rozin and Royzman 2001), the weaker-link will tend to dominate several aspects of their thinking. These aspects include how they direct their attention, their affective response to considering possible actions, and their perceptions of how impactful those actions might be. As such, we propose that decision makers intuitively evaluate possible interventions in conjunctive risk by heuristically relying on the change in the weaker-link.

While this psychology warrants further study, it seems to fit well into the category of heuristics. Classic definitions of heuristics emphasize decision makers' tendency to rely on one or relatively few key inputs to handle complex decisions (Tversky and Kahneman 1974, Gigerenzer and Goldstein 1996, Kahneman and Frederick 2002, Gigerenzer and Gaissmaier 2011), including key features of a problem (Gigerenzer and Goldstein 1996) and their initial affective responses (Slovic et al. 2007). In this vein, we argue that decision makers heuristically evaluate improvements in conjunctive risk primarily based on the impact on the weaker-link. According to this view, decision makers should use the worst-first heuristic not only when managing the complexity of conjunctive risk, but in any complex situation in which the repercussions of each possible event depend on other events. We would expect it to arise to a lesser extent in simpler domains in which the repercussions of each possible event are independent.

**Figure 8.** Size of the Bias Toward Maximizing the Weaker-Link (Negative Numbers Show a Bias Toward Maximizing the Stronger-Link) Split by Whether Participants Were Excluded for Inattentiveness Across All Studies

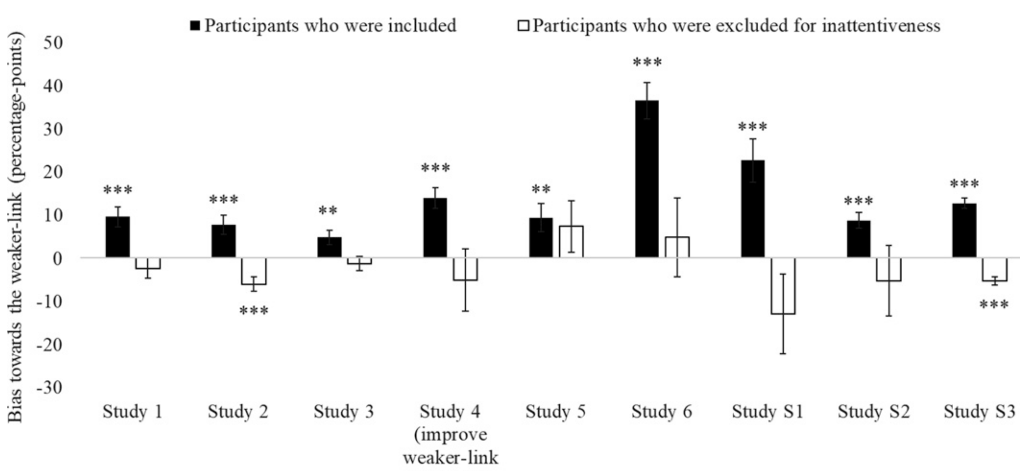


*Notes.* This chart shows the effect size of the bias toward maximizing the weaker-link split by whether participants were excluded for inattentiveness for Studies 1–6 and S1–S3. The bars for Studies 1, 3, and S3 display the coefficients of the improve-weaker-link condition variable from the main reported regression models for those studies; the bars for Study 2 display the coefficient of the improve-weaker-link condition variable as the sole predictor in a regression model; the bars for the Study 4 improve-weaker-link condition, Study 5, Study 6, and Study S1 show the difference between the proportion of participants maximizing the weaker-link and 50%; and the bars for Study S2 show the difference between the proportion of participants maximizing the weaker-link and the proportion maximizing the stronger-link. The error bars display standard errors for Studies 4, 5, 6, S1, and S2, and clustered standard errors in Studies 1, 2, 3, and S3.

$^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$.

**Figure 9.** Size of the Bias Toward Maximizing the Weaker-Link (Negative Numbers Show a Bias Toward Maximizing the Stronger-Link) Without Excluding Participants for Miscomprehension or Time Spent

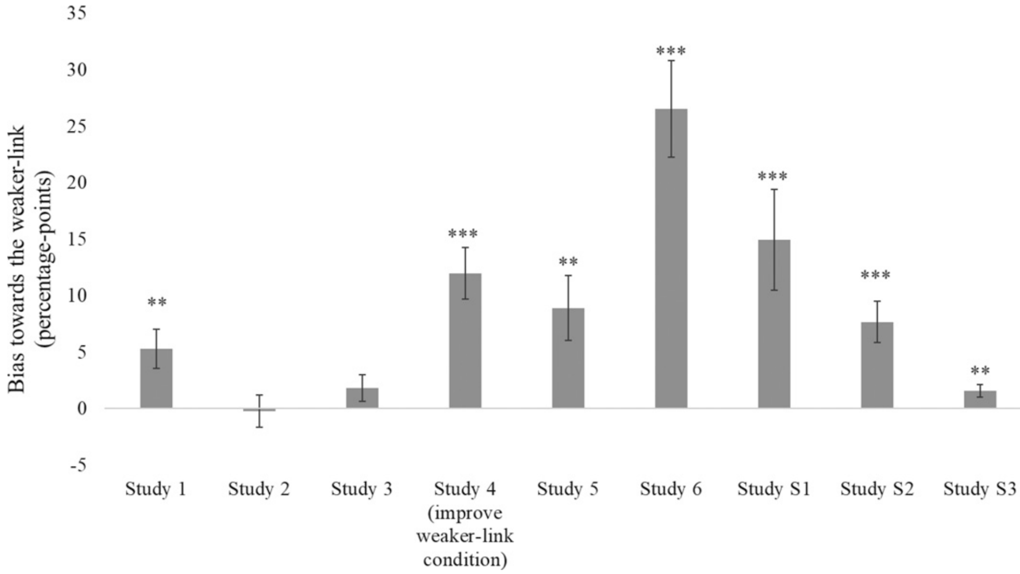


*Notes.* This chart shows the effect size of the bias toward maximizing the weaker-link without exclusions based on failing comprehension questions or time spent on questions for Studies 1–6 and S1–S3. The bars for Studies 1, 3, and S3 display the coefficients of the improve-weaker-link condition variable from the main reported regression models for those studies (but without inattentiveness exclusions); the bar for Study 2 displays the coefficient of the improve-weaker-link condition variable as the sole predictor in a regression model; the bars for the Study 4 improve-weaker-link condition, Study 5, Study 6, and Study S1 show the difference between the proportion of participants maximizing the weaker-link and 50%; and the bar for Study S2 shows the difference between the proportion of participants maximizing the weaker-link and the proportion maximizing the stronger-link. The error bars display standard errors for Studies 4, 5, 6, S1. and S2, and clustered standard errors in Studies 1, 2, 3, and S3.
$^*p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$.

Empirically, we find that the tendency to prioritize weaker-links is robust to economic incentives to preferentially improve stronger-links (Study 2), to awareness of the size of the improvement in the aggregate chance of success (Study 3 and Study S2), to dealing with reductions in chances of success rather than improvements (Study 5), and to a managerial sample facing ambiguous probabilities in an ecologically valid setting (Study 6). We also identify two boundary conditions. First, the bias toward prioritizing weaker-links is attenuated when choosing between conjunctive lotteries rather than deciding which link to improve (Study 4). Second, the bias toward prioritizing weaker-links holds particularly for attentive participants and, if anything, reverses among inattentive participants (see the previous section). Whereas we find that the worst-first heuristic is the dominant tendency overall, inattentive individuals may fall prey to a more-naive cognitive bias: simply evaluating these improvements according to their absolute magnitude.

## Alternative Accounts and Limitations

Whereas the empirical tendency toward improving weaker-links when managing conjunctive risk is clear in our data, it is worth considering alternative explanations for it. Several accounts are relatively trivial to rule out. First, any alternative explanation based on miscomprehension is ruled out by the fact that the effect is stronger in attentive participants (and is eliminated among individuals failing any comprehension checks). Second, any account that depends solely on probability weighting cannot explain why the bias is stronger when deciding between improvements in weaker-links or stronger-links than when choosing between conjunctive lotteries (Study 4) and why the effects are stronger in attentive participants (all studies). Third, any theory of how decision makers assess conjunctive risk in general, such as the configural weighted average model (Nilsson et al. 2009; see Online Supplement 12), cannot explain the results in Study 4. Finally, anticipated regret (e.g., Loomes and Sugden 1982) also seems an unlikely explanation because, no matter whether the decision maker is considering an improvement in a stronger-link or an equivalent improvement in a weaker-link, the a priori likelihood of any decision being regrettable is the same.[6]

Some alternative explanations, while unable to explain all our results, are nonetheless harder to rule out definitively or may be related to worst-first thinking. One such account is naive diversification—the idea that, when decision makers allocate investments between different options, they usually aim to have equal percentages of their portfolio in each (Benartzi and Thaler 2001). In an analogous way, when managing conjunctive risk, decision makers might aim to have an equal percentage chance of meeting each requirement. Another such alternative account is inequity aversion (Yaari and Bar-Hillel 1984, Hoffman and Spitzer 1985) with respect to the two requirements' probabilities. Just as decision makers dislike inequality across individuals, they may dislike inequality even across probabilities. Both these accounts operate on a balancing tendency, which could push decision makers toward improving weaker-links. However, the processes behind these accounts do not apply readily to our settings. For example, naive diversification has been typically studied in zero-sum portfolio settings in which the sum of all allocations is (definitionally) 100% and the diversification of independent risks is one of the explicit goals. Neither is a feature of our setting. Analogously, the concept of inequity aversion has been studied in contexts in which there are actual persons or parties to experience the unfairness of inequity. Of all our studies, only Study 6 has this feature. We nonetheless remain open to the possibility that inequity aversion and naive diversification might share some related or reinforcing psychology that could contribute to the focus on the biggest obstacle to success when managing conjunctive risk.

It is also important to consider the generalizability of our results. On the one hand, there are some reasons to think the psychology will generalize quite well. Within studies, we get the strongest results among attentive participants, and, across studies, we get the strongest result in Study 6, in which we use a sample of managers and give them a managerial scenario. In all but Study 6, judgments or decisions were incentivized. On the other hand, all of our results were laboratory experiments with relatively low stakes, so we consider the generalizability to various real-world settings to be an open question for future work. With respect to the levels of risk themselves, although we do some stimulus sampling of probabilities and potential improvements in them, and use a reasonably wide range of paradigms, there is still room to test the boundaries of the worst-first heuristic in terms of these parameters. For example, future work could explore cases in which the possible improvement in the weaker-link is smaller than the possible improvement in the stronger-link to an extreme degree, or situations in which decision makers are attempting to avoid an undesirable outcome as opposed to achieving a desirable outcome.

## Practical Implications and Conclusion

Conjunctive risk exists in a vast array of settings, any in which a decision maker requires more than one probabilistic event to achieve success or avoid failure. In managerial contexts, most projects require the successful execution of multiple, distinct components, the completion of which are uncertain. The worst-first heuristic may have important implications for managing risk in any of these settings. For example, with respect to innovative ventures, an entrepreneur must successfully navigate both technological risk (whether the technological innovation is functionally achieved) and market risk (whether sufficient demand for the innovation materializes). Our work suggests that when one of these risks is noticeably worse than the other, the manager may undervalue opportunities to improve the less-severe risk. They may strive to overcome one hurdle, only to run into the next.

Another implication involves potential inconsistencies between decisions made at different stages of a project or at different levels of management. For example, when senior executives decide which projects to start, but project managers decide how to direct resources toward completing the necessary aspects of each ongoing project, the decisions of those project managers might differ substantially from the initial expectations of the senior executives (as in Adner and Levinthal 2004). The project managers may overinvest in addressing a project's biggest challenge (i.e., the weaker-link) and neglect other critical aspects of the project. Likewise, when forced to allocate budget cuts, project managers might continue to overprioritize the project's biggest challenge.

Finally, our work raises interesting questions about the survival of individuals, firms, and societies. At the time of this writing, the world is still undergoing a global pandemic, which highlights the importance of these questions on a societal level. For humanity to survive far into the future, we must overcome multiple threats, including from pandemics, climate change, nuclear conflict, asteroids, and misaligned artificial intelligence (Bostrom 2002, Russell and Bohannon 2015, Millett and Snyder-Beattie 2017, Snyder-Beattie et al. 2019). In other words, human society's survival is subject to conjunctive risk. How to address such risks has become a major research area in global policy (Bostrom 2013, Cotton-Barratt et al. 2020, Ord 2020, Sears 2020). If policymakers' decisions are influenced by a worst-first heuristic, then they may underinvest in mitigating ostensibly less-severe threats that are more tractable. Exploring the implications of the worst-first heuristic in global risk management, as well as in managerial decision making, presents an exciting avenue for future research.

## Acknowledgment

The authors thank Akarsh Garg, Clara Han, Joy Lee, and Beren Sabuncu for research assistance.

## Appendix

**Table A.1.** Table of Contents of the Online Supplement



**Table A.2.** Links to Preregistrations

| Study | Link |
|---|---|
| Study 1 | https://aspredicted.org/n4m77.pdf |
| Study 2 | https://aspredicted.org/fy4ed.pdf |
| Study 3 | https://aspredicted.org/gb9iz.pdf |
| Study 4 | https://aspredicted.org/af7mr.pdf |
| Study 5 | https://aspredicted.org/uy6zq.pdf |
| Study 6 | https://aspredicted.org/ym6a9.pdf |
| Study S1 | https://aspredicted.org/i2i9c.pdf |
| Study S2 | https://aspredicted.org/5ci7p.pdf |
| Study S3 | https://aspredicted.org/fj96e.pdf |

## Endnotes

[1] Put differently, the half-as-large magnitude of the weaker-link improvement is exactly offset by the fact it can be pivotal twice as often; the 15-percentage-point weaker-link improvement can matter whenever the 40% chance stronger-link requirement is met, whereas the 30-percentage-point improvement in the stronger-link can matter only when the 20% chance weaker-link requirement is met. In the example, this mathematical logic is why the resulting aggregate chances of success from the two improvements are equivalent despite the improvements having different magnitudes: 35% × 40% = 20% × 70% = 14%.

[2] We formalize the distinction between managing and choosing between conjunctive risks in a mathematical model of the worst-first heuristic in Online Supplement 12, which we compare with the configural weighted average model (Nilsson et al. 2009, Jenny et al. 2014).

[3] We deviated slightly from Study 1's preregistered exclusion rules because of a typographical error affecting a comprehension check (see Online Supplement 2). This deviation does not affect significance at 0.001.

[4] In Studies 1–6 and S2, because of rounding errors, the weaker-link improvements and corresponding stronger-link improvements did not always yield identical aggregate chances of success. However, these rounding errors do not explain the bias toward weaker-links in any of these studies; see Online Supplement 15.

[5] In the example for Studies 1–3 and S3, we described a situation in which Requirement 1 had a 45% chance of being met that could be improved to 55% and Requirement 2 had a 50% chance of being met. Because the average of the possible probabilities for Requirement 1 was the same as the probability for Requirement 2, we hoped neither would be perceived as the "weaker-link" or "stronger-link," so that the comprehension question would not bias participants' responses. However, since the probability of Requirement 1 was *initially* the weaker-link, we cannot rule out this possibility completely. Fortunately, Studies 4–6 and S1–S2 do not have this problem because the example in these studies involved both weaker-link and stronger-link improvements, and the bias toward prioritizing weaker-links remained.

[6] The "regrettable" decisions in our studies would be either effortfully acquiring an improvement that does not affect the outcome or failing to acquire an improvement that would have affected the outcome. See Online Supplement 13 for further discussion of regret.

## References


Adner R, Feiler D (2019) Interdependence, perception, and investment choices: An experimental approach to decision making in innovation ecosystems. *Organ. Sci.* 30(1):109–125.

Adner R, Levinthal DA (2004) What is not a real option: Considering boundaries for the application of real options to business strategy. *Acad. Management Rev.* 29(1):74–85.

Bar-Hillel M (1973) On the subjective probability of compound events. *Organ. Behav. Human Performance* 9(3):396–406.

Benartzi S, Thaler RH (2001) Naive diversification strategies in defined contribution saving plans. *Amer. Econom. Rev.* 91(1):79–98.

Bostrom N (2002) Existential risks: analyzing human extinction scenarios and related hazards. *J. Evolution Tech.* 9, https://www.jetpress.org/volume9/risks.html.

Bostrom N (2013) Existential risk prevention as global priority. *Global Policy* 4(1):15–31.

Budescu DV, Fischer I (2001) The same but different: an empirical investigation of the reducibility principle. *J. Behav. Decision Making* 14(3):187–206.

Burson KA, Larrick RP, Lynch JG (2009) Six of one, half dozen of the other: Expanding and contracting numerical dimensions produces preference reversals. *Psych. Sci.* 20(9):1074–1078.

Cotton-Barratt O, Daniel M, Sandberg A (2020) Defence in depth against human extinction: Prevention, response, resilience, and why they all matter. *Global Policy* 11(3):271–282.

DellaVigna S, Pope D (2018) What motivates effort? Evidence and expert forecasts. *Rev. Econom. Stud.* 85(2):1029–1069.

Fan Y, Budescu DV, Diecidue E (2018) Decisions with compound lotteries. *Decision* 6(2):109.

Finucane ML, Alhakami A, Slovic P, Johnson SM (2000) The affect heuristic in judgments of risks and benefits. *J. Behav. Decision Making* 13(1):1–17.

Frederick S, Levis A, Malliaris S, Meyer A (2018) Valuing bets and hedges: Implications for the construct of risk preference. *Judgment Decision Making* 13(6):501–508.

Gigerenzer G, Gaissmaier W (2011) Heuristic decision making. *Annual Rev. Psych.* 62(1):451–482.

Gigerenzer G, Goldstein DG (1996) Reasoning the fast and frugal way: Models of bounded rationality. *Psych. Rev.* 103(4):650.

Gonzalez R, Wu G (1999) On the shape of the probability weighting function. *Cognitive Psych.* 38(1):129–166.

Hoffman E, Spitzer ML (1985) Entitlements, rights, and fairness: An experimental examination of subjects' concepts of distributive justice. *J. Legal Stud.* 14(2):259–297.

Jenny MA, Rieskamp J, Nilsson H (2014) Inferring conjunctive probabilities from noisy samples: Evidence for the configural weighted average model. *J. Experiment. Psych. Learning Memory Cognition* 40(1):203–217.

Juslin P, Nilsson H, Winman A (2009) Probability theory, not the very guide of life. *Psych. Rev.* 116(4):856–874.

Kahneman D (2011) *Thinking, Fast and Slow* (Farrar, Straus and Giroux, New York).

Kahneman D, Frederick S (2002) Representativeness revisited: Attribute substitution in intuitive judgment. Gilovich T, Griffin D, Kahneman D, eds. *Heuristics and Biases: The Psychology of Intuitive Judgment* (Cambridge University Press, New York), 49–81.

Kahneman D, Tversky A (1979) Prospect theory: An analysis of decision under risk. *Econometrica* 47(2):263–291.

Lewis J, Simmons JP (2020) Prospective outcome bias: Incurring (unnecessary) costs to achieve outcomes that are already likely. *J. Experiment. Psych. General* 149(5):870–888.

Loewenstein GF, Weber EU, Hsee CK, Welch N (2001) Risk as feelings. *Psych. Bull.* 127(2):267–286.

Loomes G, Sugden R (1982) Regret theory: An alternative theory of rational choice under uncertainty. *Econom. J.* 92(368):805–824.

Markle AB, Rottenstreich Y (2018) Simultaneous preferences for hedging and doubling down: Focal prospects, background positions, and nonconsequentialist conceptualizations of uncertainty. *Management Sci.* 64(12):5946–5959.

Millett P, Snyder-Beattie A (2017) Existential risk and cost-effective biosecurity. *Health Security* 15(4):373–383.

Nilsson H, Rieskamp J, Jenny M (2013) Exploring the overestimation of conjunctive probabilities. *Frontiers Psych.* 4:101.

Nilsson H, Winman A, Juslin P, Hansson G (2009) Linda is not a bearded lady: Configural weighting and adding as the cause of extension errors. *J. Experiment. Psych. General* 138(4):517–534.

Ord T (2020) *The Precipice: Existential Risk and the Future of Humanity* (Hachette Books, New York).

Poltrock SE, Schwartz DR (1984) Comparative judgments of multidigit numbers. *J. Experiment. Psych. Learning Memory Cognition* 10(1):32–45.

Rozin P, Royzman EB (2001) Negativity bias, negativity dominance, and contagion. *Personality Soc. Psych. Rev.* 5(4):296–320.

Russell S, Bohannon J (2015) Artificial intelligence. Fears of an AI pioneer. *Science* 349(6245):252.

Sears NA (2020) Existential security: Toward a security framework for the survival of humanity. *Global Policy* 11(2):255–266.

Shue K, Townsend RR (2021) Can the market multiply and divide? Non-proportional thinking in financial markets. *J. Finance* 76(5): 2307–2357.

Simmons JP, Nelson LD (2006) Intuitive confidence: Choosing between intuitive and nonintuitive alternatives. *J. Experiment. Psych. General* 135(3):409–428.

Slovic P (1972) *From Shakespeare to Simon: Speculations and Some Evidence about Man's Ability to Process Information* (Oregon Research Institute, Eugene, OR).

Slovic P, Lichtenstein S, Edwards W (1965) Boredom-induced changes in preferences among bets. *Amer. J. Psych.* 78(2):208–217.

Slovic P, Finucane ML, Peters E, MacGregor DG (2007) The affect heuristic. *Eur. J. Oper. Res.* 177(3):1333–1352.

Snyder-Beattie AE, Ord T, Bonsall MB (2019) An upper bound for the background rate of human extinction. *Sci. Rep.* 9(1):11054.

Tversky A, Kahneman D (1974) Judgment under uncertainty: Heuristics and biases. *Science* 185(4157):1124–1131.

Weber EU, Johnson EJ (2009) Mindful judgment and decision making. *Annual Rev. Psych.* 60(1):53–85.

Yaari ME, Bar-Hillel M (1984) On dividing justly. *Soc. Choice Welfare* 1(1):1–24.