Working paper, last updated 4/19/2023

Biases in Improvement Decisions:

People Focus on the Relative Reduction in Bad Outcomes

William H. Ryan*

Haas School of Business

University of California, Berkeley

wryan@berkeley.edu

Stephen M. Baum*

Haas School of Business

University of California, Berkeley

stephen_baum@berkeley.edu

Ellen R. K. Evers

Haas School of Business

University of California, Berkeley

erkevers@berkeley.edu

* Denotes co-first authorship

Abstract

People often decide whether to invest scarce resources—such as time, money, or energy—to improve the chances of a positive outcome. For example, a doctor might decide whether to utilize scarce medicine to improve a patient’s chances of recovery, or a student might decide whether to study a few additional hours to increase their chance of passing an exam. Eleven studies ($N = 5,382$) find evidence that people behave as if they focus on the relative reduction in bad outcomes caused by such improvements. As a consequence, the same improvements (e.g., 10 percentage point improvements) are valued very differently depending on whether one’s initial chances of success are high or low. This focus on the relative reduction of bad outcomes drives risk preferences that violate normative standards (Studies 1a-2a), is amplified when decisions become more consequential (Study 2b), and leads even experienced professionals to make suboptimal decisions (Study 3).

Keywords: risk, improvement, decision-making, bias, regret

Statement of Relevance

People commonly encounter risk in their daily lives. Decisions about one’s finances, health, and occupation all entail accepting and managing risks. While psychologists have long been interested in how people approach such risks, little research has explored how people think about and value *improving* the chances that a risk produces a positive outcome. This is puzzling, as many risky decisions are improvement decisions (e.g., a student deciding whether to expend additional time studying to improve the chances that they pass an exam). The current research finds that people make these improvement decisions by considering how they would feel if they

48 were to obtain a *bad* outcome. This leads to suboptimal lay and expert decision-making in a
49 variety of consequential contexts (e.g., medicine and finance) and strongly implies that risk
50 preferences do not reflect a stable personality trait, but are dictated by context-specific
51 psychological forces.

52

Many decisions involve choosing whether to invest money, time, or effort to improve the chances of obtaining some positive outcome. Executives determine how much budget to allocate to campaigns designed to boost sales, students decide how many hours they will spend studying to improve their chances of passing an exam, and intelligence analysts evaluate costly new initiatives which might reduce the likelihood of terror attacks. In a world in which money, time, and effort are all limited, such improvement decisions are consequential: Investing resources in one improvement often means neglecting others.

While these improvement decisions are common, there is relatively little research on how people make them. Of course, a large literature in psychology, economics, and public policy has examined how people think about and value risky prospects (e.g., people's willingness-to-pay for a 30% chance of winning a reward; see Mata et al., 2018). In turn, it seems straightforward to assume that people's valuations of an improvement (e.g., going from a 30% to 40% chance of winning) might be roughly equivalent to the difference in their valuations of two individual risky prospects (e.g., the difference in valuations between a 30% versus 40% chance of winning a reward; see Wu & Gonzalez, 1996). However, a recent paper (Lewis & Simmons, 2020) shows that people often prefer to improve their chances of success when those chances are already initially high. The authors argue that this is because people consider their expected final outcome in addition to the improvement itself (a *prospective outcome bias*).

This manuscript proposes a psychology of how people make improvement decisions. Our expectations are informed by the literature on regret (e.g., Zeelenberg, 1999; Loomes & Sugden, 1982). Regret arises when people compare a realized outcome to a superior counterfactual alternative which they could have obtained had they acted differently (Bleichrodt & Wakker, 2015; Loomes & Sugden, 1982). Because experiencing regret is unpleasant, people work to

minimize the regret they expect to feel for their actions (Zeelenberg, 1999). We propose that people evaluate improvements by focusing on how they would feel if they were to obtain a bad outcome, and how much they could have done to prevent that bad outcome from occurring.¹

To illustrate, consider Mike, a medical student preparing for an exam. While Mike currently has a 10% chance of passing the exam, staying up all night will increase his chances of passing to 20%. We believe that when making this improvement decision, Mike will consider how he would feel if he *didn't* study and subsequently failed the exam. And, since staying up all night would not have substantially reduced his chances of failing (the 10 percentage point improvement translates into a 1/9th reduction in the proportion of failing outcomes), Mike would believe that he could not have done much to prevent failure. In turn, Mike would minimally value staying up all night .

Conversely, consider Kate, another student preparing for the same exam. Although Kate can also stay up all night to increase her chances of passing an exam by 10 percentage points, she currently has an 80% chance of passing the exam. If Kate imagines her feelings if she forgoes studying and fails, she will likely believe that she could easily have prevented failure, as the 10 percentage point improvement would lead to a full ½ reduction in her proportion of failing outcomes. Thus, we suggest that Kate will greatly value spending the night studying. While Mike and Kate both face the same improvement decision (pulling an all-nighter to increase their chances of passing by 10 percentage points), we expect them to treat these improvements very differently.

More formally, we suggest that people naturally consider the relative reduction in bad outcomes caused by an improvement without fully considering how initially likely bad outcomes

¹ We examine situations with two possible outcomes (i.e., win or not win, pass or fail) and label the inferior outcome the *bad outcome*.

are to occur. Therefore, the same 10 percentage point improvement might be perceived as extremely impactful or relatively minor depending on one's initial chances of success (see Figure 1).²

Figure 1

Bad Outcomes Eliminated by the Same Improvement with Different Initial Chances of Success

A. Improving an 80% chance of success to a 90% chance:

Beginning with an 80% chance of success...

...going from 80% to 90% improves 10% of all outcomes...

...and fully eliminates 50% of bad outcomes

B. Improving a 10% chance of success to a 20% chance:

Beginning with a 10% chance of success...

...going from 10% to 20% improves 10% of all outcomes...

...but only eliminates 12% of bad outcomes

Moreover, this focus on the proportion of bad outcomes eliminated should yield a distinctive pattern of valuations of improvements. Because the same improvement reduces an ever-greater proportion of bad outcomes as initial probabilities of success increase, we hypothesized that a convex pattern of valuations of improvements would emerge across the full range of initial probabilities (see Figure 2).

If people indeed value improvements in this manner, then they will likely make suboptimal decisions. Specifically, decision-makers may prefer improving risky prospects with initially higher chances of success (e.g., improving a 90% chance of a reward to 95%) over more dramatically improving risky prospects with initially lower chances of success (e.g., improving a 10% chance of a reward to 20%). Our suggestion that the biased anticipation of regret drives

² We deviate from past formal models of regret that assume decision-makers minimize regret by equally considering all possible outcomes (e.g. Bell, 1982; Loomes & Sugden 1982). We instead propose that decision-makers consider potential regret conditional on obtaining a bad outcome.

114 **Figure 2**115 *Proportional Reductions in Bad Outcomes for Ten Percentage Point Improvements*

116

117 *Note.* Illustrates the proportion of bad outcomes eliminated by a 10 percentage point118 improvement as initial probability of success increases.³

119 these suboptimal decisions is theoretically and empirically meaningful. For example, many

120 decision-making biases arise from a lack of deliberation (Kahneman & Frederick, 2002) or

121 motivation (Petty et al., 1981). Therefore, it is generally assumed that making decisions more

122 consequential—for example, by increasing reward magnitudes—increases decision quality

123 (Thaler, 1986). However, since regret arises when people obtain an inferior outcome and blame

124 themselves (Zeelenberg & Pieters, 2007), increasing the gap between an inferior, bad outcome

³ Supplementary Materials section IV shows this pattern does not follow from expected utility theory, prospect theory or regret theory.

and superior, good outcome may *amplify* anticipated regret. Our studies test this regret-based logic directly.

Eleven studies investigate how people think about and value improvements. Studies 1a-1g use diverse methods and measures to test if participants' valuations of improvements follow the hypothesized convex pattern Figure 2 depicts. Studies 2a-2b explore anticipated regret's underlying role in these valuations. Finally, Study 3 examines how lay and expert decision-makers make choices about improvements across varied domains. First, we report a pilot study testing if people treat improvements differently than static risky outcomes.

Open Practices Statement

The Supplementary Materials detail additional analyses and exploratory measures, as well as analyses including participants who failed preregistered exclusion criteria. All studies—except the Pilot Study and Study 1a—were preregistered. All data, materials, preregistrations, and analysis scripts have been made publicly available at:

https://osf.io/acvrm/?view_only=139ea21a8fc14dbbbca9248246286945.

Pilot Study

In the pilot study, 97 participants indicated both their willingness-to-pay to enter \$10 lotteries (*enter* condition), and their willingness-to-pay to change their chances of winning \$10 lotteries (*change* condition). Figure 3 compares participants' explicit valuations of improvements (e.g., the average valuation of changing a 10% chance of winning \$10 to a 20% chance; purple line) with their implied valuations of improvements (e.g., the difference between the average valuation of a 10% chance of winning \$10 and the average valuation of a 20% chance of winning \$10; green line).

Figure 3 makes clear how improvements are treated very differently than (mathematically

Figure 3

WTP for Ten Percentage Point Increase in Chances of Winning \$10

Note. Participants' mean willingness-to-pay for explicit (purple line) and implied (green line)

improvements in the Pilot Study. Error bars are +/- 1 SE.

equivalent) static risky prospects. The flat green line reveals that participants' implied valuation of each improvement was roughly similar across all initial probabilities. That is, the gap between participants' valuations of two individual risks did not change across the full spectrum of initial probabilities. However, the purple line highlights how participants' valuations of explicit improvements varied greatly based on the initial probability of success. Further, the slope of the purple line in Figure 3 suggests that valuations of improvements follow a convex pattern. Next, we turn to studies designed to test for this pattern more thoroughly.

Study 1

Study 1 consists of seven different studies. Each study consisted of a series of trials in which participants indicated their valuation of an improvement (e.g., “How much are you willing to pay to change a 50% chance of winning \$10 to a 60% chance of winning \$10?”). While all seven studies followed this same overarching structure, each study used slightly different stimuli, methods, and measures. For example, some studies were hypothetical, while others were fully incentive-compatible (via the BDM procedure; Becker et al., 1964). Some studies asked about a limited set of initial probabilities (e.g., initial probabilities from 10% to 90%, in 10 percentage point increments), while others were more expansive (asking about 32 improvements, with many initially non-round probabilities). And while Studies 1a-1g asked participants to indicate their willingness-to-pay for improvements, Study 1f instead asked participants to indicate their willingness-to-work for improvements. We predicted that regardless of the elicitation method, number of trials, or specific set of probabilities sampled, a convex pattern of valuations of improvements would emerge.

Method

Participants and design. Depending on the study, participants completed between 16 and 43 trials. In each trial, participants had an initial probability of winning \$5 or \$10 (this varied across studies). Participants indicated how much they were willing to pay (or how many captchas they were willing to solve) to increase their chance of winning by ten percentage points (p.p). Study 1c also measured valuations of 5 p.p. increases. Results for these trials are reported in the Supplemental Materials; these trials followed a similar convex pattern as the 10 p.p. increases. Table 1 provides additional information on the design of each study.

182 **Table 1**183 *Design and features of Studies 1a-1g.*

Study	N	Prize	# Trials	Sample	Description	vs. Classic Regret Theory Model*	vs. Prospective Outcome Bias Model**
1a	109	\$10	16	AMT	Finds convex pattern of valuations	$\chi^2(1) = 395.91, p < .001$	>.99
1b	101	\$10	43	AMT	Measured 5 p.p. and 10 p.p. improvements	$\chi^2(1) = 903.11, p < .001$	>.99
1c	127	\$10	22	AMT	Includes more trials with higher initial probabilities	$\chi^2(1) = 617.71, p < .001$	>.99
1d†	96	\$5	20	Students	Fully incentive compatible through in-person BDM procedure	$\chi^2(1) = 48.93, p < .001$	>.99
1e†	203	\$5	25	AMT	Larger sample size and incentive compatible	$\chi^2(1) = 2,246.80, p < .001$	>.99
1f†	109	\$5	32	Students	Expanded set of trials	$\chi^2(1) = 620.24, p < .001$	>.99
1g†	117	\$5	25	AMT	Participants indicate willingness to work for improvements	$\chi^2(1) = 1,181.80, p < .001$	>.99

184 † Incentive compatible study

185 * The conditional regret and classic regret theory models were compared using a Chi-square test, which was possible because the two models are nested. A

186 significant Chi-square test means that the conditional regret model fits the data significantly better than the classic regret theory model.

187 ** Because the conditional regret and prospective outcome bias models were not nested, they could not be compared with a Chi-square test. Instead, a

188 comparison based on AIC was used. Here, >.99 means that there is a more than 99% probability that the conditional regret model is correct given the comparison

189 model and underlying data. Additional details are given in the Supplementary Materials.

We were concerned that participants might misunderstand the study instructions and give their valuation of the final outcome as opposed to the improvement. To prevent this misunderstanding, we used extremely strict (preregistered) exclusion criteria. In all of Studies 1a-1g, participants completed two attention check trials where the improvement gave no benefit (i.e., it was zero p.p.). Participants who failed this check were excluded. Each study also included comprehension questions which asked participants to display an understanding of the consequences of their choices. Analyses including all respondents—which are consistent with what we report in the main text—can be found in the Supplementary Materials.

Results

Figure 4 presents participants' mean willingness-to-pay for 10 percentage point improvements. The initial probability of each lottery is plotted on the x-axis. For example, the values at 60% represent participants' mean willingness-to-pay to increase their chances of winning \$10 from 60% to 70%. As hypothesized, when changes in chances begin to reduce the proportion of bad outcomes more dramatically, participants' valuations of such chances rapidly increase. To test our hypothesis more formally, we compared a model which contained a predictor for the relative reduction in bad outcomes (which we term the conditional regret model) against models that follow from alternative theories (e.g., models that follow from regret theory, Loomes & Sugden, 1982, or prospective outcome bias, Lewis & Simmons, 2020). In all of 1a-1g, the conditional regret model best fit the data (see Table 1). Full information on robustness checks—as well as additional model comparisons—can be found in the Supplemental Materials.

Moreover, the convex pattern depicted in Figure 4 is particularly notable when one considers that Studies 1a-1g have a fully within-subjects design. Specifically, many participants display very different valuations of the same (10 p.p.) improvements. For example, in Study 1a,

Figure 4

Mean WTP for a Ten Percentage Point Improvement in Chance to Win

Note. Participants' mean willingness-to-pay for a 10 percentage point improvement by initial probability (Studies 1a-1f), and participants' mean willingness-to-work (i.e., solve Captchas) for a 10 percentage point improvement by initial probability (Study 1g). Error bars are ± 1 SE.

all trials improved participants' chances of winning \$5 by 10 p.p. (i.e., an expected value of \$0.50). However, participants behaved *risk-averse* below initial probabilities of 70% (e.g., participants were willing to pay an average of \$0.15 to change a 10% chance of winning \$5 to a 20% chance), and *risk-seeking* above initial probabilities of 70% (e.g., those same participants were willing to pay an average of \$0.61 to change an 80% chance of winning \$5 to a 90% chance). In other words, the same participants—within the same context and using the same measures—displayed both an affinity for and an aversion to identical risky improvements. This

behavior is hard to reconcile with the common belief that risk attitudes are a stable individual difference (e.g., Von Neumann & Morgenstern, 1944).

Study 2

Studies 1a-1g utilized an array of methods and measures to find that valuations of improvements follow a convex pattern. We have proposed that when evaluating improvements, people: (1) naturally focus on bad outcomes, and (2) anticipate how much regret they would feel if they decided to forgo that improvement. Our next two studies more directly test these propositions. Study 2a manipulates what people focus on when evaluating improvements, while Study 2b indirectly manipulates anticipated regret by varying the consequences of a bad outcome.

Study 2a

If people naturally focus on bad outcomes when considering improvements, a manipulation which draws attention to bad outcomes should affect valuations of improvements only minimally. In contrast, a manipulation which draws attention to good outcomes should attenuate—or even reverse—the bias we have documented thus far. Study 2a tests this prediction via a framing manipulation.

Method

Participants and design. 2,421 Americans were recruited from Prolific Academic in exchange for payment. 91 participants who failed a preregistered attention check were excluded from analyses. Participants were randomly assigned to one of six between-subjects conditions in a 2(initial probability: 20% or 80%) X 3(frame: neutral, proportional increase, proportional decrease) full-factorial design. In the 20% [80%] initial probability condition, participants read that they had a 20% [80%] chance of winning \$10.00 and that they had the opportunity to increase their chances of winning by 10 percentage points (see Figure 5 for a stylized depiction

of the instructions). In all conditions, participants used a \$0 to \$3.00 slider to indicate their willingness-to-pay to increase their winning chances. In the neutral condition, participants were presented with this improvement depicted solely in terms of percentage points. In the proportional increase condition, participants read that a 10 percentage point improvement would increase their chances of winning by $1/8$ (in the 80% condition) or $1/2$ (in the 20% condition); in the proportional decrease condition, participants read that a 10 percentage point improvement would decrease their chances of losing by $1/2$ (in the 80% condition) or $1/8$ (in the 20% condition).

Figure 5

Information Presented by Condition, Study 2a

Results

We submitted participants' willingness-to-pay to a 2(initial probability: 20% or 80%) $\times$ 3(frame: neutral, proportional increase, proportional decrease) ANOVA. This revealed a main effect of initial probability, $F(1, 2324) = 20.05, p < .001, \eta_p^2 = .009$. In other words, participants

valued changing an 80% chance of winning \$10 to a 90% chance more than they valued changing a 20% chance of winning \$10 to a 30% chance. However, as Figure 6 reveals, this pattern varied across the three framing conditions. A significant Initial Probability X Frame interaction confirmed this, $F(2, 2324) = 4.34, p = .013, \eta_p^2 = .004$.

In the neutral condition, participants were willing to pay more to change an 80% chance of winning \$10 to a 90% chance ($M = \$0.90, SD = \0.75) than to change a 20% chance of winning \$10 to a 30% chance ($M = \$0.74, SD = \0.66), $t(804) = 3.17, p = .002, d = 0.22$. This is consistent with the results of Studies 1a-1d. This pattern also appeared in the proportional decrease condition. When improvements were framed in terms of the relative reduction in losing outcomes, participants valued changing the 80% chance to a 90% chance ($M = \$1.01, SD = \0.79) more than they valued changing the 20% chance to a 20% chance ($M = \$0.78, SD = \0.69), $t(762) = 4.15, p < .001, d = 0.30$. Crucially, there was no significant interaction between the neutral and proportional decrease conditions ($F < 1$).

However, as the rightmost bars in Figure 6 reveal, the proportional increase frame produced a different pattern of results. In this condition, there was no difference in valuations between the 20% ($M = \$0.86, SD = \0.69) and 80% initial probability gambles ($M = \$0.87, SD = \$0.67, t < 1$). This attenuation was confirmed by a significant interaction between the neutral and proportional increase conditions, $F(1, 1562) = 4.21, p = .040, \eta_p^2 = .003$. Unsurprisingly, there was also a significant interaction effect between the proportional increase and proportional decrease conditions, $F(1, 1520) = 8.26, p = .004, \eta_p^2 = .005$. In sum, while a manipulation which prompted participants to consider the relative reduction in bad outcomes did not significantly alter the effect of initial probability on valuations of improvements, a parallel manipulation which instead prompted participants to consider good outcomes eliminated this effect. This

pattern is consistent with our suggestion that people evaluate improvements by focusing on the relative reduction in bad outcomes. **Figure 6**

Mean WTP to Improve Chances by Ten Percentage Points

Note. Participants' mean willingness-to-pay to change their chances of winning a given lottery by 10 percentage points in Study 2a. Error bars are ± 1 SE.

Study 2b

Study 2b was designed to provide more direct evidence that the anticipation of regret underlies valuations of improvements. People anticipate regret when they feel responsible for obtaining an inferior outcome (Zeelenberg & Pieters, 2007). Whereas some biases reduce in size when decisions become more consequential (Kahneman & Frederick, 2002; see also Sloman, 1996), the anticipation of regret increases when decisions grow in importance (Frijda et al., 1989). Therefore, if anticipated regret underlies valuations of improvements, making a bad outcome *worse*—and thereby increasing the gap between an anticipated bad outcome and a

superior counterfactual alternative—should increase the size of the observed effect. Study 2b tests this using a paradigm which directly manipulates the severity of a bad outcome.

Two additional elements of Study 2b’s design are worth noting. First, while our previous studies presented participants with opportunities to increase the likelihood of a good outcome, Study 2b instead presented participants with an opportunity to decrease the likelihood of a bad outcome. We expected our core effect to be robust to this change. Second, Study 2b provided a conservative test of our hypothesis: it tests if participants will value a smaller improvement (i.e., a 9 percentage point improvement) over a larger one (i.e., a 10 percentage point improvement) if that smaller improvement leads to a greater relative reduction in bad outcomes.

Method

Participants and design. Participants were 1,607 Americans recruited from AMT. We excluded 19 participants who failed a preregistered attention check. All participants were randomly assigned to one of four between-subjects conditions in a 2(initial probability: 60% or 10%) X 2(consequences: mild or severe) full-factorial design.

Participants began by reading that they had the opportunity to purchase a pill which would reduce their risk of getting a seasonal cold. Our first manipulation varied participants’ initial probability of catching the cold. In the 60% [10%] initial probability condition, participants read that they had a 60% [10%] chance of catching the cold, and that if they purchased and consumed the pill, they would instead have a 50% [1%] chance of getting the cold. A second manipulation varied the consequences of catching the cold. The cold’s symptoms were described to participants as being mild (and lasting one week) or severe (and lasting three weeks). After reading about the cold, the likelihood of catching it, and the pill (see Figure 7 for a stylized depiction of the instructions), participants provided their willingness-to-pay for the pill.

Willingness-to-pay was elicited on a slider scale ranging from \$0.00 to \$100.00 (defaulting to \$0.00 and taken to two decimal places). Participants also completed a series of exploratory measures which asked how much regret they would anticipate experiencing if they did not purchase the pill and did catch the cold. We discuss these measures in complete detail in the Supplementary Materials.

Figure 7

Information Presented by Condition, Study 2b

Results

We submitted participants' willingness-to-pay for the pill to a 2(initial probability: 60% chance or 10% chance) X 2(consequences: severe or mild) ANOVA. This revealed an expected main effect of Consequences, $F(1, 1584) = 57.51, p < .001, \eta_p^2 = .035$, indicating that participants' willingness-to-pay was higher when the cold carried more severe symptoms. Furthermore, there was a main effect of Initial Probability, $F(1, 1584) = 82.89, p < .001, \eta_p^2 = .050$; participants valued the pill more when it eliminated a larger proportion of bad outcomes. Crucially, these two main effects were qualified by the predicted Consequences X Initial

Probability interaction, $F(1, 1584) = 12.88, p < .001, \eta_p^2 = .008$ (see Figure 8). When the cold's symptoms were mild, participants were willing to pay more for the pill when it reduced their chances of getting sick from 10% to 1% ($M = \$14.60, SD = \16.45) than they were when it reduced their chances of getting sick from 60% to 50% ($M = \$9.71, SD = \14.29), $t(794) = 4.48, p < .001, d = 0.32$. This difference was amplified when the consequences of getting sick were more severe: Here, participants' willingness-to-pay for a pill which reduced their chances of getting sick from 10% to 1% ($M = \$24.84, SD = \22.77) was greater than their willingness to pay for a pill that reduced their chances of getting sick from 60% to 50% ($M = \$13.59, SD = \16.21), $t(790) = 8.07, p < .001, d = 0.57$.

In sum, Study 2b found that increasing the severity of a bad outcome increased the extent to which participants were sensitive to the relative reduction in bad outcomes. In addition, exploratory analyses—detailed in the Supplementary Materials—highlight how Study 2b's measures of anticipated regret followed a similar interaction pattern as the willingness-to-pay measures in Figure 8. Taken together, these patterns are consistent with our suggestion that participants value improvements by considering the extent to which they would regret realizing a bad outcome.

Study 3

Study 2b found that people value smaller improvements more than larger improvements when the former lead to a greater relative reduction in bad outcomes. This implies that people may *choose* objectively inferior improvements over better alternatives. Such a bias would have broad implications for high-stakes decision-making in domains ranging from medicine (e.g., doctors deciding which patients to prioritize treating) to commerce (e.g., managers deciding which project to invest in). To this point, Study 3 examines how both expert and non-expert

Figure 8

Mean WTP for a Pill to Reduce Chances of Catching a Cold

Note. Mean willingness-to-pay for a pill which would decrease participants' chances of catching a cold in Study 2b. Error bars are ± 1 SE.

decision-makers make such choices in different naturalistic contexts. For example, Study 3 presented all participants with improvement decisions in a medical context (e.g., choosing which of two patients would receive a treatment increasing their chances of survival), and included both a lay sample and a sample of medical professionals (nurses, doctors, etc.). This allowed us to directly compare how lay participants and experienced professionals make improvement decisions within a given domain. *A priori*, we found three different predictions to be plausible.

The first prediction is straightforward: Typically, experience is thought to improve decision-making (Kahneman & Klein, 2009). Experienced decision-makers are often more

motivated and deliberate more than novice decision-makers (Ericsson et al., 1993). This perspective suggests that expert professionals—within their domain of expertise—will make more optimal improvement decisions than lay participants.

However, a second plausible outcome is that that expertise will not alleviate the bias we have seen thus far. Expertise is typically acquired via feedback based on repeated experience. Crucially, decisions to (not) improve risky prospects do not always allow for learning. While, in prospect, outcomes are probabilistic (e.g., an 80% chance of passing an exam), experiences are dichotomous (either a student passes an exam or does not). A student who decides to stay up all night studying to improve their chances of passing will likely not be able to pinpoint if that extra effort actually made a difference in their score. For this reason, the primary channel for learning—repeated exposure to feedback that allows for an audit of the quality of one’s decision—often does not exist for many improvement decisions. This reasoning suggests that having significant experience may not guard against suboptimal improvement decisions.

Finally, a third line of reasoning suggests that experience may actually *exacerbate* the bias we have seen thus far. It is entirely believable that experienced professionals see decisions in their domain of expertise as more consequential than non-experts do. Study 2b found evidence that making decisions more consequential actually increased participants’ biased valuation of improvements. In turn, professionals might make suboptimal decisions about improvements at a greater rate than non-experts.

Method

Procedure and design. Study 3 consisted of three waves. In each wave, we recruited a separate sample of participants. We recruited 198 medical professionals (e.g., nurses/doctors), 198 gig workers (e.g., Uber/Doordash drivers), and 301 lay participants. The medical

professionals and gig workers were recruited by Qualtrics Panels, and the lay participants through Prolific Academic. We excluded 192 participants who failed preregistered comprehension checks (113 medical professionals, 30 gig workers, and 49 lay participants). Across all three waves of data collection, all participants made choices in seven different domains (see Table 2); in each domain, participants completed four choice trials. The order in which participants saw the domains was randomized, and within each domain, the order of the four trials was randomized. For example, medical professionals made four choices in the medical domain; lay participants and gig workers made the same four choices. We repeated this procedure across all seven domains. This feature of the design allowed us to directly compare experts' and non-experts' choices in a given domain.

Each trial presented participants with two options. Both options had some initial probability of success, and participants had to choose which option to improve. Within each domain, there was one control trial, and three different experimental trials.

Table 2

Domains and Choice Scenarios Presented to Participants, Study 3

Domain	Choice scenario
Medicine	Giving a treatment to one of two critically ill patients
Gig economy	Focusing on meeting one of two bonus goals
Public health	Distributing a vaccine for one of two different diseases
Business	Allocating budget to one of two different projects
Law	Doing additional preparation for one of two upcoming court cases
Education	Spending time studying for one of two upcoming exams
Finance	Receiving additional tickets for one of two different raffles

In the control trials, the size of the improvement was always greater for one option than for the other option. For example, a control trial in the medical domain presented participants with two patients, both of whom initially had a 10% chance of survival. Participants had to select which patient they would prioritize treating. If they selected the first patient, that patient's chance of survival would improve from 10% to 20%; if they selected the second patient, that patient's chances of survival would improve from 10% to 19% (see Branch A of Figure 9). In this way, all of the control trials presented participants with a clearly superior improvement - going from 10% to 20% is both a larger absolute increase in chances than the alternative (10 percentage points vs. nine percentage points) and leads to a larger relative reduction in bad outcomes (11% eliminated vs. 10% eliminated). Thus, the control trials serve as a benchmark. We expected the vast majority of participants to prefer the superior improvement to the inferior alternative.

The experimental trials had a slightly different structure. Just as with the control trials, participants saw two options and had to choose which option to improve. However, unlike the control trials, one option in the experimental trials always corresponded with a larger absolute increase in chances, while the other option led to a larger relative reduction in bad outcomes (see Branch B of Figure 9). This detail allows us to test if participants' sensitivity to the relative reduction in bad outcomes might lead to suboptimal choices. If medical decisions are made with an eye towards maximizing the number of lives saved, improving a patient's chances of survival from 10% to 20% is—all else equal—a more optimal decision than improving a patient's chance of survival from 80% to 89% (a 10 percentage point vs. nine percentage point improvement). Yet, we expected that many participants in these experimental trials would choose the smaller improvements—e.g., 80% to 89%—which correspond with a greater relative reduction in bad outcomes (11% eliminated vs. 45% eliminated).

There were three types of experimental trials. The three types of trial varied whether the smaller improvement was: either nine percentage points, seven percentage points, or five percentage point (see Branches B-D of Figure 9). This feature allowed us to test if participants would prefer improvements that led to larger relative reductions in bad outcomes even when these improvements were vastly inferior (e.g., five percentage points vs. 10 percentage points).

Figure 9

Information Presented to Participants by Condition, Study 3

At the end of the survey, all participants completed two comprehension check questions asking them to correctly identify the consequences of their choices, as well as several professional and demographic questions (see the Supplementary Materials for complete details).

Results

Results for all participants across the entire range of domains and trials are shown in Figure 10. For the control trials, there was a clear (and unsurprising) pattern; participants overwhelmingly preferred the superior improvements. That is, only ~5% of participants preferred the smaller to larger improvements (see purple bars in Figure 10). However, in the

experimental trials, participants made suboptimal decisions: 56% of participants chose the smaller improvement in the nine p.p. trials (yellow bars in Figure 10). To more formally test this, we conducted a linear mixed-effects model predicting choice of the smaller improvement from the trial (either the control or the nine p.p. experimental trial) with random intercepts by participant. This model revealed that participants were significantly more likely to choose the lower absolute value option in the nine p.p. condition relative to the control condition ($\beta = .52$, $t(6,564) = 60.00$, $p < .001$), an effect which remained robust with alternate specifications (see the Supplementary Materials).

An examination of the green and yellow bars in Figure 10 showcases how participants made suboptimal decisions across the five, seven, and nine p.p. trials at a (roughly) similar frequency. In the seven p.p. trials, 52% of participants chose the suboptimal improvement, while in the five p.p. trials, 49% did. In other words, approximately half the time, participants preferred improvements that were twice as small, but led to a greater relative reduction in bad outcomes. For example, 45% of medical professionals chose a treatment which would improve a patient's chances of survival by five percentage points over a treatment which would improve a patient's chances of survival by 10 percentage points. The fact that so many professionals were willing to choose an option which was only half as impactful as the alternative is a testament to how aversive anticipated regret can be.

We next turn to more fine-grained tests of the effects of expertise on improvement decisions. Here, we conducted two analyses. The first compared gig workers' choices within the gig work trials against lay participant and medical professionals' choices within those same trials (see Figure 11A), and the second compared medical professionals' choices within the medical

476 **Figure 10**477 *Choice of the Smaller Improvement Option in Study 3 by Sample, Scenario, and Trial Condition*

478

479 *Note.* The percentage of participants choosing an option which gave a smaller absolute improvement (i.e., five, seven, or nine p.p.)

480 over an option which gave a larger ten p.p. improvement.

Figure 11*Choice of the Smaller Improvement in Gig and Medical Domains by Sample*

Note. Choices of the suboptimal option giving the smaller improvement in the gig domain for all three samples (A, left) and in the medical domain for all three samples (B, right).

trials against lay participant and gig workers' choices within those trials (see Figure 11B).

For gig workers, we found no evidence that expertise led to better or worse decision-making (i.e., 76% of gig workers chose the smaller improvement in nine p.p. condition, as compared to 78% for medical professionals and 77% for lay participants; *ns* on all comparisons). However, in the medical sample, we actually found that medical professionals were *more* likely to make suboptimal choices (48% choosing smaller improvement in nine p.p. condition vs. 32% for gig workers and 31% for lay participants; $\beta = .17$, $t(503) = 3.09$, $p < .001$), an effect which held even with a variety of robustness checks (see the Supplemental Materials).

This increase might have occurred for a variety of reasons. Experts might see decisions within their field as being highly important, and thus be more likely to anticipate feeling regret when making such decisions (as implied by Study 2b). Moreover, the experience of regret is purely dichotomous; once outcomes are realized, experts either experience regret or do not.

Experts—with months, years, or a lifetime of making decisions in a specific domain—might be most quick to recall situations in which their initial chances of success were high, they did not improve their chances, and they subsequently realized a bad outcome. This biased recall might lead improvement attached to initially high probabilities to be particularly attractive. Independent of each of these (speculative) interpretations of our results, our data strongly suggests that expertise alone may not be sufficient to attenuate the biased valuation of improvements.

General Discussion

Many decisions require people to consider whether to expend scarce resources to improve the likelihood of obtaining a good outcome. However, little prior research has examined how these improvement decisions are made. Eleven studies demonstrate that people make these decisions by focusing on the relative reduction in bad outcomes. Consistent with our theorizing, participants' valuations of improvements followed a convex pattern. This pattern appeared for a range of probabilities and payoffs, for investments of effort or money, and in hypothetical and incentive compatible settings. Our final study found that participants' sensitivity to the relative reduction in bad outcomes led to suboptimal choices in consequential domains.

Our proposal that the phenomenology of these improvement decisions can be characterized by the (biased) anticipation of regret has important implications. Mathematical models of regret—which do not assume that people primarily focus on bad outcomes—are commonly used to model behavior in a variety of domains (e.g., Loomes & Sugden, 1982). The Supplementary Materials sketch out how the most prominent of these models can be updated to better capture valuations of improvements. Further, many decision-making biases are depicted as occurring from a lack of motivation or deliberation. As decisions become more consequential—and the motivation to deliberate increases—the size of these biases shrink (Kahneman &

Frederick, 2002). However, because regret emerges when people feel responsible for a bad outcome (Loomes & Sugden, 1982; Zeelenberg, 1999), making decisions more consequential can actually *strengthen* regret-based biases. The worse a bad outcome, the more responsibility—and thus regret—people feel (see Frijda et al., 1989). The results of Study 2b are consistent with this.

Our suggestion that anticipated regret guides decisions about risky improvements also helps build a more cumulative understanding of how emotions dictate risk preferences. Researchers have long proposed that risk preferences are a “personality trait”: those who are risk-averse or risk-seeking in one domain will also be risk-averse or risk-seeking in another (Blais & Weber, 2006, p. 33). However, evidence for this general risk appetite is limited (see Fox & Tannenbaum, 2011; but see also Frey et al. 2017). Moreover, recent findings suggest that risk preferences may not even be stable *within* domains. For example, participants who pay more for a risk-increasing bet also pay more for a risk-decreasing hedge of that bet (Frederick et al., 2018). Our model’s proposal that decision-makers focus on bad outcomes provides an intuitive psychological explanation for Frederick et al.’s findings (see Supplementary Materials) and rationalizes people displaying both risk-seeking and risk-aversion (see Figure 4). Our findings cannot be easily accommodated by risk-as-a-personality-trait perspectives (e.g., Blais & Weber, 2006), but are consistent with the suggestion that anticipated emotions often dictate risky decisions (e.g., Zeelenberg et al., 1996).

While our manuscript focuses on how people consider improving *single* risky prospects, we believe a similar psychology governs choices *between* multiple risky prospects. Consider a decision-maker who is choosing between an 80% chance of winning \$10 or a 90% chance of winning \$8. While the first option’s expected value is greater, choosing the second option yields

half as many losing outcomes, and thus is preferred per our theory's logic. This pattern, if empirically supported, might help to explain puzzles in the decision-making literature, including why people prefer lower expected value options over alternatives (i.e., the Ellsberg paradox), have highly inconsistent preferences for the same options (i.e., transitivity cycles), and value illusory certainty (i.e., pseudocertainty effects, see Supplemental Study 3).

Finally, as the bias we have identified persists even when professionals make consequential decisions, it is important to pinpoint interventions to reduce it. One possible intervention involves reframing probabilistic improvements (e.g., going from a 10% to a 20% chance of a patient surviving) in terms of the average expected outcome (e.g., increasing the average survival rate from 10 to 20 patients out of 100). Anticipating the experience of regret is easy in individual decisions with concrete implications for one person. But, when making decisions with an average effect on a much larger group, attention drifts away from any individual outcome—good or bad—and towards an *aggregate* set of outcomes. Indeed, Supplemental Study 4 finds that reframing improvements decisions in this way attenuates the bias we observe (23% of participants choose the suboptimal improvement with reframing, 35% without; $\beta = 0.57$, $t(400) = 2.54$, $p = .01$). This result is akin to the identifiable victim effect, whereby an individuated person's suffering provokes a stronger emotional response than a larger group's collective plight (Jenni & Loewenstein, 1997).

Conclusion

Understanding how decision-makers think about risk has long been of interest to psychologists, organizations, and governments. Many dominant theories of risky decision-making do not explicitly distinguish between how people think about risk and how people think about risky *improvements*. Our findings suggest that decision-makers disproportionately focus on

567 the relative reduction in bad outcomes when considering improvements. Thus, as the initial
568 probability of realizing a good outcome increases, participants behave as though it is more likely
569 that they will experience a negative emotion. We hope future work directly investigates how this
570 psychology might characterize decision-making in applied spheres such as insurance, medicine,
571 and goal progress.
572

References

- 573
574 Becker, G. M., Degroot, M. H., & Marschak, J. (1964). Measuring utility by a single-response
575 sequential method. *Behavioral Science*, 9(3), 226–232.
576 <https://doi.org/10.1002/bs.3830090304>
- 577 Bell, D. E. (1982). Regret in Decision Making under Uncertainty. *Operations Research*, 30(5),
578 961–981. <https://doi.org/10.1287/opre.30.5.961>
- 579 Blais, A.-R., & Weber, E. U. (2006). A Domain-Specific Risk-Taking (DOSPERT) scale for
580 adult populations. *Judgment and Decision Making*, 1(1), 33–47.
581 <https://doi.org/10.1017/S1930297500000334>
- 582 Bleichrodt, H., & Wakker, P. P. (2015). Regret Theory: A Bold Alternative to the Alternatives.
583 *The Economic Journal*, 125(583), 493–532. <https://doi.org/10.1111/econj.12200>
- 584 Connolly, T., & Zeelenberg, M. (2002). Regret in Decision Making. *Current Directions in*
585 *Psychological Science*, 11(6), 212–216. <https://doi.org/10.1111/1467-8721.00203>
- 586 Ericsson, K. A., Krampe, R. T., & Tesch-Römer, C. (1993). The role of deliberate practice in the
587 acquisition of expert performance. *Psychological Review*, 100(3), 363–406.
588 <https://doi.org/10.1037/0033-295X.100.3.363>
- 589 Fox, C. R., & Tannenbaum, D. (2011). The Elusive Search for Stable Risk Preferences. *Frontiers*
590 *in Psychology*, 2. <https://doi.org/10.3389/fpsyg.2011.00298>
- 591 Frey, R., Pedroni, A., Mata, R., Rieskamp, J., & Hertwig, R. (2017). Risk preference shares the
592 psychometric structure of major psychological traits. *Science Advances*, 3(10), e1701381.
593 <https://doi.org/10.1126/sciadv.1701381>
- 594 Frijda, N. H., Kuipers, P., & ter Schure, E. (1989). Relations among emotion, appraisal, and

- 595 emotional action readiness. *Journal of Personality and Social Psychology*, 57(2), 212–
596 228. <https://doi.org/10.1037/0022-3514.57.2.212>
- 597 Frederick, S., Levis, A., Malliaris, S., & Meyer, A. (2018). Valuing bets and hedges:
598 Implications for the construct of risk preferences. *Judgment and Decision Making*, 13(6),
599 501-508.
- 600 Jenni, K. E., & Loewenstein, G. (1997). Explaining the “Identifiable Victim Effect.” *Journal of*
601 *Risk and Uncertainty*, 14, 235-257. <http://doi.org/10.1023/A:1007740225484>
- 602 Kahneman, D., & Frederick, S. (2002). Representativeness Revisited: Attribute Substitution in
603 Intuitive Judgment. In T. Gilovich, D. Griffin, & D. Kahneman (Eds.), *Heuristics and*
604 *Biases* (1st ed., pp. 49–81). Cambridge University Press.
605 <https://doi.org/10.1017/CBO9780511808098.004>
- 606 Kahneman, D., & Klein, G. (2009). Conditions for intuitive expertise: A failure to disagree.
607 *American Psychologist*, 64(6), 515–526. <https://doi.org/10.1037/a0016755>
- 608 Lewis, J., & Simmons, J. P. (2020). Prospective outcome bias: Incurring (unnecessary) costs to
609 achieve outcomes that are already likely. *Journal of Experimental Psychology: General*,
610 149(5), 870–888. <https://doi.org/10.1037/xge0000686>
- 611 Loomes, G., & Sugden, R. (1982). Regret Theory: An Alternative Theory of Rational Choice
612 Under Uncertainty. *The Economic Journal*, 92(368), 805.
613 <https://doi.org/10.2307/2232669>
- 614 Mata, R., Frey, R., Richter, D., Schupp, J., & Hertwig, R. (2018). Risk Preference: A View from
615 Psychology. *Journal of Economic Perspectives*, 32(2), 155–172.
616 <https://doi.org/10.1257/jep.32.2.155>
- 617 Petty, R. E., Cacioppo, J. T., & Goldman, R. (1981). Personal involvement as a determinant of

- 618 argument-based persuasion. *Journal of Personality and Social Psychology*, 41(5), 847–
619 855. <https://doi.org/10.1037/0022-3514.41.5.847>
- 620 Sloman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin*,
621 119(1), 3–22. <https://doi.org/10.1037/0033-2909.119.1.3>
- 622 Thaler, R. H. (1986). The Psychology and Economics Conference Handbook: Commentary on
623 papers by H. Simon, H. Einhorn and R. Hogarth, and A. Tversky and D. Kahneman,
624 *Journal of Business*.
- 625 Von Neumann, J., & Morgenstern, O. (1944). Theory of games and economic behavior. In
626 *Theory of Games and Economic Behavior*. Princeton University Press.
- 627 Wu, G., & Gonzalez, R. (1996). Curvature of the Probability Weighting Function. *Management*
628 *Science*, 42(12), 1676–1690. <https://doi.org/10.1287/mnsc.42.12.1676>
- 629 Zeelenberg, M. (1999). Anticipated regret, expected feedback and behavioral decision making.
630 *Journal of Behavioral Decision Making*, 12(2), 93–106.
631 [https://doi.org/10.1002/\(SICI\)1099-0771\(199906\)12:2<93::AID-BDM311>3.0.CO;2-S](https://doi.org/10.1002/(SICI)1099-0771(199906)12:2<93::AID-BDM311>3.0.CO;2-S)
- 632 Zeelenberg, M., Beattie, J., van der Pligt, J., & de Vries, N. K. (1996). Consequences of Regret
633 Aversion: Effects of Expected Feedback on Risky Decision Making. *Organizational*
634 *Behavior and Human Decision Processes*, 65(2), 148–158.
635 <https://doi.org/10.1006/obhd.1996.0013>
- 636 Zeelenberg, M., & Pieters, R. (2007). A Theory of Regret Regulation 1.0. *Journal of Consumer*
637 *Psychology*, 17(1), 3–18. https://doi.org/10.1207/s15327663jcp1701_3