

THE PSYCHOLOGY OF IMPROVEMENTS

Joshua Lewis

A DISSERTATION

in

Operations, Information and Decisions

For the Graduate Group in Managerial Science and Applied Economics

Presented to the Faculties of the University of Pennsylvania

in

Partial Fulfillment of the Requirements for the
Degree of Doctor of Philosophy

2020

Supervisor of Dissertation

Joseph Simmons
Professor of Operations, Information and Decisions

Graduate Group Chairperson

Nancy Zhang, Ge Li and Ning Zhao Professor, Professor of Statistics

Dissertation Committee:

Joseph Simmons, Professor of Operations, Information and Decisions
Alice Moon, Assistant Professor of Operations, Information and Decisions
Deborah Small, Laura and John J. Pomerantz Professor of Marketing

ACKNOWLEDGMENT

I appreciate the financial support from the Operations, Information, and Decisions Department, the Wharton Behavioral Laboratory, the Wharton Risk Center Ackoff Doctoral Student Fellowship, and the Baker Retailing Center Graduate Student Grant.

I also want to thank my committee members, Alice Moon, Deborah Small, and Joseph Simmons, for giving me excellent feedback. An extra thanks goes to Joseph Simmons for being an incredible advisor and an inspiration. None of this would be possible without him.

ABSTRACT

THE PSYCHOLOGY OF IMPROVEMENTS

Joshua Lewis

Joseph P. Simmons

How do people decide whether to incur costs to increase their likelihood of success? In investigating this question, we offer a theory called *prospective outcome bias*. According to this theory, people tend to make decisions that they expect to feel good about after the outcome has been realized. Because people expect to feel best about decisions that are followed by successes—even when the decisions did not *cause* those successes—they will pay more to increase their chances of success when success is already likely (e.g., people will pay more to increase their probability of success from 80% to 90% than from 10% to 20%). We find evidence for prospective outcome bias in nine experiments. In Study 1, we establish that people evaluate costly decisions that precede successes more favorably than costly decisions that precede failures, even when the decisions did not cause the outcome. Study 2 establishes, in an incentive-compatible laboratory setting, that people are more motivated to increase higher chances of success. Studies 3–5 generalize the effect to other contexts and decisions and Studies 6–8 indicate that prospective outcome bias causes it (rather than regret aversion, waste aversion, goals-as-reference-points, probability weighting, or loss aversion). Finally, in Study 9, we find evidence for another prediction of prospective outcome bias: people prefer small increases in the probability of large rewards (e.g., a 1% improvement in their chances of winning \$100) to large increases in

the probability of small rewards (e.g., a 10% improvement in their chances of winning \$10).

TABLE OF CONTENTS

ACKNOWLEDGMENT	II
ABSTRACT.....	III
LIST OF TABLES	VI
LIST OF FIGURES	VII
THE PSYCHOLOGY OF IMPROVEMENTS.....	1
ENDNOTES	59
APPENDIX.....	65

LIST OF TABLES

Table 1. Stimuli Used in Studies 2 and 6.....	14
Table 2. Scenarios Used in Studies 5 and 8.....	43
Table 3. Study 9 Stimuli	50
Table A1. Table of Contents of the Online Supplemental Material	65

LIST OF FIGURES

Figure 1. Study 1 Screenshot	15
Figure 2. Study 1 Results	18
Figure 3. Study 6 Results	29
Figure 4. Study 7 Screenshots.....	36
Figure 5. Study 7 Results	40
Figure 6. Study 8 Results	46
Figure 7. Study 9 Screenshot	51
Figure 8. Study 9 High vs. Low Results	53
Figure 9. Study 9 5% Increase vs. 25% Increase Results	54

THE PSYCHOLOGY OF IMPROVEMENTS

People often have to decide whether to incur a cost to increase their likelihood of success. For example, a salesperson must decide whether to spend time chasing down a promising lead, a patient must decide whether to undergo the pain of getting a flu vaccine, a citizen must decide whether to stand in line to vote for her preferred candidate, and an academic must decide whether to revise a manuscript one more time before submitting it for publication. Although decisions of this type are ubiquitous, we do not fully understand how people make them. In particular, we do not yet know whether people are more motivated to increase their chances of success when success is currently unlikely or when it is already likely.

In our first attempt to answer this question, we (confidently) predicted that people would be more motivated to improve their likelihood of success when their chances of succeeding were low (e.g., 10%) than when they were high (e.g., 80%). There are at least two reasons why we generated this hypothesis. First, when the chances of winning are already very high, people may more easily rationalize that there is no need for additional effort (Fishbach & Finkelstein, 2012). Second, because people often encode differences as ratios, going from a 10% to a 20% chance of success may feel like a bigger improvement than going from an 80% to a 90% chance of success (e.g., Tversky & Kahneman, 1981).

We learned very quickly—after one pilot study—that our prediction was wrong and that the opposite is true: People are more motivated to improve their chances of success when those chances are high than when they are low. In this paper, we will show you the

evidence for this effect, and we will suggest that *prospective outcome bias* is why it happens.

Prospective Outcome Bias

When people evaluate decisions after the outcomes of those decisions have been observed, they succumb to *outcome bias*: wise decisions that result in negative outcomes are judged negatively, whereas unwise decisions that result in positive outcomes are judged positively (Baron & Hershey, 1988). The tendency to evaluate the quality of decisions based on their outcomes is pervasive, as it has been demonstrated to emerge in decisions about ethics (Gino, Moore, & Bazerman, 2009; Gino, Shu, & Bazerman, 2010), business (Marshall & Mowen, 1993), the military (Lipshitz, 1989), finances (Ratner & Herbst, 2005), public safety (Tinsley, Dillon, & Cronin, 2012), sports (Lefgren, Platt, & Price, 2015), and health care (Caplan, Posner, & Cheney, 1991). Importantly, people succumb to outcome bias even when evaluating their *own* decisions, and these biased evaluations can affect their future decisions, as has been shown for investment decisions in the laboratory (Ratner & Herbst, 2005), real-life hurricane evacuation decisions (Tinsley et al., 2012), and the player selection decisions of National Basketball Association coaches (Lefgren et al., 2015).

Although research has shown how outcome bias affects *subsequent* decisions, we know of no research that has considered whether it affects the *original decisions themselves*. But, there are reasons to believe that it does. First, there is a large body of research showing that decision makers are influenced by how they expect to evaluate their

decisions in the future (Loomes & Sugden, 1982; Mellers & McGraw, 2001; van de Ven & Zeelenberg, 2011; Zeelenberg, 1999; Zeelenberg & Pieters, 2004). That is, people are more likely to make decisions that they expect to feel good about afterward than decisions that they expect to feel bad about afterward. Second, because outcome bias is so pervasive, people may be able to anticipate that they will feel better about a decision that is followed by a good outcome than about a decision that is followed by a bad outcome, *even if the decision itself does not cause that outcome*. Thus, people may be more inclined to make decisions that are expected to be followed by good outcomes than identically wise decisions that are expected to be followed by bad outcomes. Importantly, this preference may be especially strong for costly decisions (such as investments or purchases), which more often require the justification that a positive outcome would provide. We refer to this account as *prospective outcome bias*.

Predictions of Prospective Outcome Bias

If we apply prospective outcome bias to the questions that motivated this research, then we can make two novel predictions. First, we predict that people will invest more resources to improve their chances of a desired outcome when the chances are already high than when the chances are low. To understand why, imagine that someone pays a fee to increase the probability of winning a prize. It will feel less painful for her to have paid that fee if she wins the prize than if she loses the prize because it is easier to justify having paid the fee if she wins. And, because she can anticipate this, she will be more likely to pay the fee when she is likely to win the prize (and thus when she is likely to be able to easily

justify the fee) than when she is unlikely to win the prize (and thus when she is unlikely to be able to justify the fee).

In general, when people consider whether to incur costs to modestly improve *already-high* chances of success, they will expect the costs to feel relatively painless (since they expect to be successful). They will therefore be inclined to incur those costs. In contrast, when people consider whether to incur costs to modestly improve *low* chances of success, they will expect the costs to feel relatively painful (since they expect to be unsuccessful). They will therefore be *disinclined* to incur those costs. Thus, prospective outcome bias predicts that, all else equal, people will be more likely to invest resources to increase the probability of success when that probability is already high than when it is low.

It is important to note that outcome bias applies much more to *costly* decisions that are intended to *improve* one's chances of success (e.g., buying extra raffle tickets to improve one's chances of winning a prize), than to *costless* decisions that merely *affect* one's chances of success but are *not* intended to improve them (e.g., declining to buy extra raffle tickets, or even selling existing tickets). This is because people need to justify costly decisions much more than they need to justify costless decisions. For example, imagine that you had decided to buy extra tickets for a raffle, thus increasing your chances of winning it. If you lose the raffle, it is hard to justify the money you spent on those extra raffle tickets, and so you are likely to negatively evaluate the decision to buy them. Now instead imagine that you had decided *not* to buy extra tickets for a raffle. Even if you lose this raffle, it is still easy to justify the decision *not* to buy the extra tickets on

the grounds that you saved money, and so you are *unlikely* to evaluate this decision negatively. The distinction between initially costly decisions (which are more affected by outcome bias) and initially costless decisions (which are less affected by outcome bias) is important because it explains why, when their chances of success are already high, people find the decision to invest resources to further increase those chances more appealing relative to the decision *not* to invest. (Study S1 in Online Supplement 2 provides an empirical demonstration of this phenomenon.)

Prospective outcome bias also makes a second prediction: Holding constant the potential increase in the expected value of the lottery, people will be more likely to invest to acquire a *smaller* improvement in the chances of winning a *larger* prize than to acquire a *larger* improvement in the chances of winning a *smaller* prize. We expect this result because potentially unnecessary costs will feel more justified, and thus less painful, if the decision to incur those costs is accompanied by a more positive outcome (i.e., a larger prize).

These two predictions follow from a core assumption of prospective outcome bias: people's willingness to invest to improve their chances of success is *not* driven solely by the value that they place on the improvement itself. Instead, because the eventual outcome will determine whether an investment feels justified, people's willingness to invest is also driven by how much value they place on the *expected* outcome. We can very simply model people's willingness to invest as a weighted average of the *improvement* in the expected value of their chances of success $\Delta E(S)$ and the *final* expected value of their chances of success $E(S)$ (i.e., the prospective outcome). So, if we denote the weight on the expected

value of the final chances of success as w (such that $0 < w < 1$), we can define willingness to invest as¹:

$$\text{Willingness To Invest} = (1 - w) \cdot \Delta E(S) + w \cdot E(S).$$

Our key predictions directly follow from this equation.² First, people will incur greater costs to attain a given improvement in their chances of success (a constant value of $\Delta E(S)$) when their initial chances of success are higher, and thus the final expected value of these chances ($E(S)$) is also higher.³ Second, when the best possible outcome is of higher value, the final expected value of the chances of success ($E(S)$) is also higher, and so again, people are willing to incur greater costs to attain a given improvement in the expected value of their chances of success ($\Delta E(S)$).⁴

We tested the assumptions and predictions of prospective outcome bias in the nine experiments presented below. Along the way, we also rule out alternative accounts of these phenomena, including anticipated regret, goals-as-reference-points, and probability weighting.

Research Overview

In this article, we present nine studies investigating prospective outcome bias. In Study 1, we present evidence for a critical assumption of this theory: people expect to evaluate a decision that precedes a success more highly than a decision that precedes a failure, even when the decision did not cause the success or failure. We find that people do anticipate that they will judge future decisions based on their outcomes. Then, in Studies 2–9, we document an important consequence of prospective outcome bias: people are more

motivated to increase high chances of success than low chances of success. Study 2 establishes this effect in an incentive-compatible laboratory setting, Study 3 shows that it generalizes to decisions to forego a reward to maintain probabilities, Study 4 shows that it generalizes to randomly sampled probabilities, and Study 5 shows that it generalizes to a wide variety of real-world decisions. Studies 6–8 provide further evidence for the effect, while helping to establish that prospective outcome bias is the cause of it. Study 6 shows that people’s greater propensity to decide to improve high versus low chances of success emerges even when they know that they will learn the impact of that decision, a result that rules out anticipated regret and waste aversion as plausible mechanisms. Study 7 rules out probability weighting, goals-as-reference-points, and loss aversion explanations, and Study 8 shows that the more outcome bias people expect to have, the more motivated they are to improve high versus low chances of success. Finally, in Study 9, we conceptually replicate the results of Study 2, while also showing evidence for another effect that follows from prospective outcome bias: holding the potential improvement in expected value constant, people prefer small increases in the probability of a large reward to large increases in the probability of a small reward.

In all of our studies, we report all of our measures, manipulations, exclusions, and rules for determining sample size. In some studies, we used attention checks and comprehension checks as a basis for excluding participants, and those checks are presented in detail in Online Supplement 6. The full breakdowns of exactly which participants and observations were excluded from each study are presented in Online Supplement 1. All of our studies were preregistered on AsPredicted.org, and the links to those preregistrations

are in the Appendix. We perfectly followed all of our preregistered analysis plans and exclusion rules, with two exceptions. First, in Study 8, we deviated in our computation of the predictor variable for a conceptual reason that we discuss when we present that study. (Study 8's findings continue to be very significant if we conduct the analysis described in our preregistration.) Second, although we preregistered to use easier-to-interpret ordinary least squares (OLS) regressions instead of logistic regressions for analyses with binary dependent variables (Studies 2, 3, 5, 6, and 9), we have honored the review team's request to report the results of logistic regressions in these instances. All of the statistical tests that we report using binary dependent variables have the same level of significance using either model, and we report the preregistered OLS analyses in Online Supplement 7. The online supplemental materials, complete study materials, and all of our data and code are available at this website: <https://osf.io/57jku/>. The research was approved by the University of Pennsylvania's Institutional Review Board (protocol #829286).

Study 1

Our theory of prospective outcome bias hinges on the assumption that people expect to judge initially costly decisions more favorably when they are followed by good outcomes rather than bad outcomes, *even when the decisions did not influence those outcomes*. The goal of Study 1 was to directly test whether this assumption is valid.

In this study, we asked participants to imagine that they were endowed with either a 30% chance or a 60% chance of winning a raffle, and that they had purchased 10 additional raffle tickets so as to increase their win probability by 10 percentage points. All

participants also imagined that their decision to purchase the raffle tickets turned out to be inconsequential. Specifically, those endowed with a 30% chance learned that they had subsequently lost the raffle, whereas those endowed with a 60% chance learned that they had subsequently won the raffle, but would still have won it even if they had not purchased additional tickets. Participants then evaluated the purchase decision. We predicted that they would evaluate it more positively when they subsequently won the raffle than when they subsequently lost the raffle. This finding would be consistent with our hypothesis: people expect to more positively evaluate inconsequential (but costly) decisions that are followed by good outcomes than inconsequential (but costly) decisions that are followed by bad outcomes.

Method

Participants

We recruited participants from Amazon's Mechanical Turk (MTurk). Participants received \$0.30 for completing the study. We decided in advance to collect data from 500 participants. In the event of multiple responses from a single MTurk ID or IP address, we preregistered to include the original response only, resulting in 67 exclusions. After all preregistered exclusions, including some described below, our final sample comprised 403 participants (mean age = 36.1, 44.5% female).⁵

Procedure

Participants completed an online survey. The survey began by telling participants that they would have to rate the quality of a decision to buy extra raffle tickets. On the next

page, they were presented with an attention check question that they had to answer in a way that indicated that they understood this instruction. If they failed to answer correctly on their first attempt, the survey prevented them from continuing and we collected no additional data from them.

Participants who passed the attention check then read a scenario in which they imagined participating in a raffle. They were randomly assigned to either a *win* condition or a *lose* condition. The *win* [*lose*] condition scenario read as follows:

Imagine that you participate in a raffle for a \$200 Amazon gift card. The winning raffle ticket will be randomly drawn from 100 raffle tickets, numbered 1 through 100. Imagine that you originally were given numbers 1–60 [1–30], and that somebody offers you the chance to buy 10 extra numbers (61–70 [31–40]) for \$10. Imagine that you buy the extra numbers for \$10, and that the winning number turns out to be 50. As a result, you win [do not win] the \$200 gift card, but you would [would not] have won regardless of whether you decided to buy the extra tickets for \$10.

How good a decision was it to buy the extra tickets, numbered 61–70 [31–40]?

All participants evaluated the decision to buy the extra raffle tickets on a 7-point scale ranging from 1 = *extremely bad* to 7 = *extremely good*. This measure was our dependent variable. On the same page, we asked participants to justify their evaluation in a text box.

To ensure that the participants had properly understood the scenario, we then asked them two comprehension check questions. In the first, they had to indicate whether the

winning ticket was one of their own tickets, and in the second, they had to indicate whether the winning ticket was one of the extra tickets that they had bought. At the end of the survey, participants entered demographic information.

Results and Discussion

We preregistered to exclude participants who answered either of the two comprehension checks incorrectly, resulting in 28 exclusions from the *win* condition and 18 exclusions from the *lose* condition. This left us with 403 participants for our main analysis, 194 in the *win* condition and 209 in the *lose* condition. The statistical significance of our results does not change if we include all participants.

We expected participants to evaluate the inconsequential (but costly) decision to buy additional raffle tickets more favorably when the decision was followed by a good outcome than when it was followed by a bad outcome. Consistent with this hypothesis, we found that participants rated the decision much more favorably in the *win* condition ($M = 4.59$, $SD = 1.65$) than in the *lose* condition ($M = 3.18$, $SD = 1.73$), $t(401) = 8.39$, $p < .001$.

This result suggests that people will more favorably evaluate a costly decision to improve their probability of success when success is attained than when it is not, even if the decision did not cause the outcome. Thus, people may expect to judge a costly decision to improve their chances of success more positively if success is already likely than if it is unlikely.

Study 2

In Study 1, we found that people expect to feel better about costly decisions that are followed by success than those that are not, even when the decisions are not themselves pivotal. This result suggests that, if people's decisions are influenced by how they expect to feel about those decisions after the associated outcomes have been realized, then they should be more likely to pay to increase win probabilities when success is already likely than when success is unlikely. Thus, in the context of a prize draw, people should be more willing to pay to increase high win probabilities than low win probabilities. We tested this prediction in Study 2, in which participants made incentive-compatible decisions about whether to improve their chances of winning various prizes by investing effort.

Method

Participants

Participants were paid \$10 for an hour-long laboratory session in a northeastern university, and this study was a 15-minute part of this session. We preregistered to collect participants from a week-long series of lab sessions, aiming to get 200 participants in total.⁶ We also preregistered to include only the original response of any participant who completed the study more than once, but none of them did. After all preregistered exclusions, including some described below, our final sample comprised 158 participants (mean age = 20.4, 70.9% female).

Design

Participants decided whether to improve their chances of winning each of 15 prizes by typing “ab” a given amount of times (DellaVigna & Pope, 2018). We manipulated the

win probabilities across prizes, and thus within-subjects. For each prize, we held constant the probability increase that participants could attain, and we randomly assigned participants to either low potential win probabilities (i.e., below 50%) or high potential win probabilities (i.e., above 50%). For each prize, we also ensured that the potential high win probabilities and the potential low win probabilities were (on average) the same distance from 50% so as to rule out any condition differences that might be driven by people treating moderate probabilities differently from more extreme probabilities.

Table 1 lists all of the prizes, as well as the win probabilities and required amount of typing associated with each prize. For example, the first row of Table 1 shows that, on one trial, participants were asked whether they would be willing to type “ab” 110 times to increase their probability of winning a \$5 Amazon gift card either from 3% to 15% (*low probabilities* condition) or from 85% to 97% (*high probabilities* condition). The last row of Table 1 shows that, on a different trial, participants were asked whether they would be willing to type “ab” 20 times to increase their probability of winning post-its either from 3% to 17% or from 83% to 97%. As can be seen from these examples, the size of the probability increase and the extremity of the probabilities were held constant across conditions.

Table 1
Stimuli Used in Studies 2 and 6

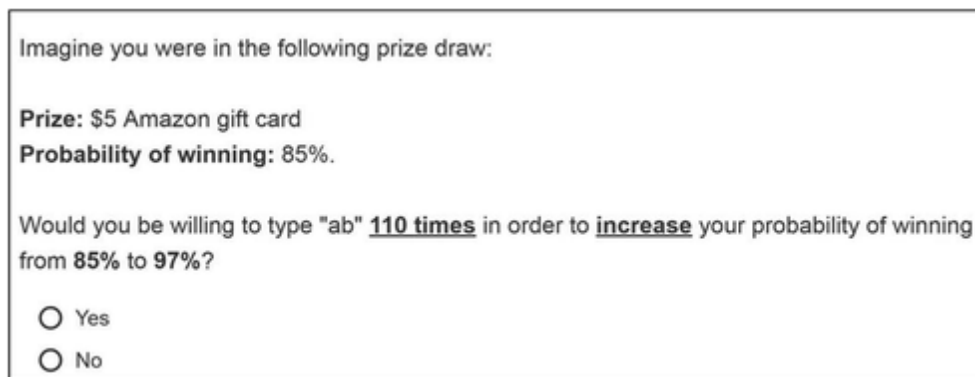
Prize	Required instances of typing “ab” to increase win probability	Possible increase in win probability	
		Low probabilities	High probabilities
Amazon \$5 gift card	110	From 3% to 15%	From 85% to 97%
Starbucks \$5 gift card	105	From 5% to 28%	From 72% to 95%
Dunkin’ donuts \$5 gift card	85	From 4% to 16%	From 84% to 96%
Pack of oreos	75	From 3% to 9%	From 91% to 97%
Granola bar	65	From 4% to 11%	From 89% to 96%
CVS \$5 gift card	90	From 4% to 19%	From 81% to 96%
Flashlight	65	From 5% to 26%	From 74% to 95%
Mini-bag of M&Ms	70	From 3% to 13%	From 87% to 97%
Four Hershey’s kisses	60	From 5% to 20%	From 80% to 95%
Notebook	55	From 4% to 16%	From 84% to 96%
McDonald’s \$5 gift card	80	From 3% to 17%	From 83% to 97%
Bottle of coke	70	From 3% to 6%	From 94% to 97%
Highlighter	45	From 5% to 23%	From 77% to 95%
Pen	30	From 5% to 16%	From 84% to 95%
Post-its	20	From 3% to 17%	From 83% to 97%

Procedure

At the beginning of the survey, we explained to participants that they would decide whether to improve their chances of winning each of 15 prizes by typing “ab” on a keyboard a given amount of times. To ensure that participants took their decisions seriously, we truthfully told them that we would randomly select one of these 15 prize draws to conduct for real, and we displayed a selection of the prizes on a table at the front of the room. Participants then answered two comprehension questions about these instructions. The first question required them to indicate that they were being asked to decide whether to *increase* their chances of winning a prize, as opposed to whether to *enter* into a prize draw at a given probability. The second question required them to indicate that the survey would randomly select one of the prize draws to conduct for real. To give participants some intuition for how effortful the task would be, we also asked participants to practice typing “ab” 100 times before they made any decisions. Then, on

the next page, they indicated how motivated they were to win each of the 15 prizes on a 7-point scale ranging from 1 = *extremely unmotivated* to 7 = *extremely motivated*.

Participants then made their decisions, which were presented one at a time on the computer screen and in a random order. For each decision, we informed participants of the prize, what their baseline probability of winning would be, and what their increased probability of winning would be if they agreed to type “ab” the required amount of times. Figure 1 shows what participants saw on the Amazon gift card trial.



Imagine you were in the following prize draw:

Prize: \$5 Amazon gift card
Probability of winning: 85%.

Would you be willing to type "ab" **110 times** in order to increase your probability of winning from 85% to 97%?

☐ Yes
☐ No

Figure 1. Screenshot of the *high probabilities* condition for the Amazon gift card trial in Study 2.

After they made their decisions, the survey randomly selected one of the prize draws to conduct for real, and informed participants of which prize draw had been selected. If the participant previously *agreed* to type “ab” the required amount of times for this prize, they were required to do so before they could complete the survey, and then they found out whether they won at the *increased* probability. If the participant previously *declined* to type “ab” the required amount of times for the selected prize, they completed the survey and then immediately found out whether they won at the *baseline* probability (i.e., without doing any extra typing). When participants won a prize, they were asked to raise their hand

so a research assistant could make sure they received their prize. At the end of the survey, participants entered their demographic information. On their way out of the laboratory, a research assistant gave participants their \$10 participation fee along with any prize that they won.

Results and Discussion

To ensure that the participants included in our main analyses understood the instructions and that the decisions included in our analyses pertained to prizes that participants were sufficiently motivated to win, we preregistered two additional exclusion rules. First, we excluded data from any participant who failed the first comprehension question on their first attempt. Second, we excluded decisions about prizes for which participants rated their motivation to win at less than a 4 (*neutral*) on a 7-point scale (1 = *extremely unmotivated* to 7 = *extremely motivated*).⁷ Note that both of these exclusion rules were based on participants' responses to questions that were asked *before* they were assigned to the *high* versus *low probabilities* condition for any specific item, and thus could not have been influenced by their condition assignments. This left us with 158 participants and 1,262 observations for our main analysis.

Each participant provided up to 15 observations of the dependent variable, one for each prize for which they rated their motivation to win as at least “neutral”. Our key dependent variable captured whether, for each prize, the participant decided to type “ab” the required amount of times to increase the probability of winning (1 = *yes*, 0 = *no*). To test the hypothesis that people are more likely to invest effort to increase high rather than

low win probabilities, we used a logit model to regress this variable on a *high probabilities* condition dummy variable (1 = *high probabilities* condition; 0 = *low probabilities* condition). We included fixed effects for each prize, and we accounted for the nonindependence of observations by clustering standard errors by participant.

Figure 2 displays the results for each item as well as collapsed across items. Consistent with our hypothesis, we found that participants were significantly more likely to invest effort to improve high win probabilities than low win probabilities, $b = 1.24$, clustered $SE = .17$, $p < .001$, $OR = 3.45$ (95% CI [2.49, 4.79]). This finding is consistent with prospective outcome bias. At high probabilities of winning, participants know that because they are likely to win, their investment of effort is likely to seem justified; as a result, they are motivated to expend that effort to improve their chances. At low probabilities of winning, participants know that since they are unlikely to win, their investment of effort is *not* likely to seem justified; as a result, they are less motivated to expend that effort to improve their chances.

Study 2: Participants were more likely to expend effort to increase high win probabilities than low win probabilities

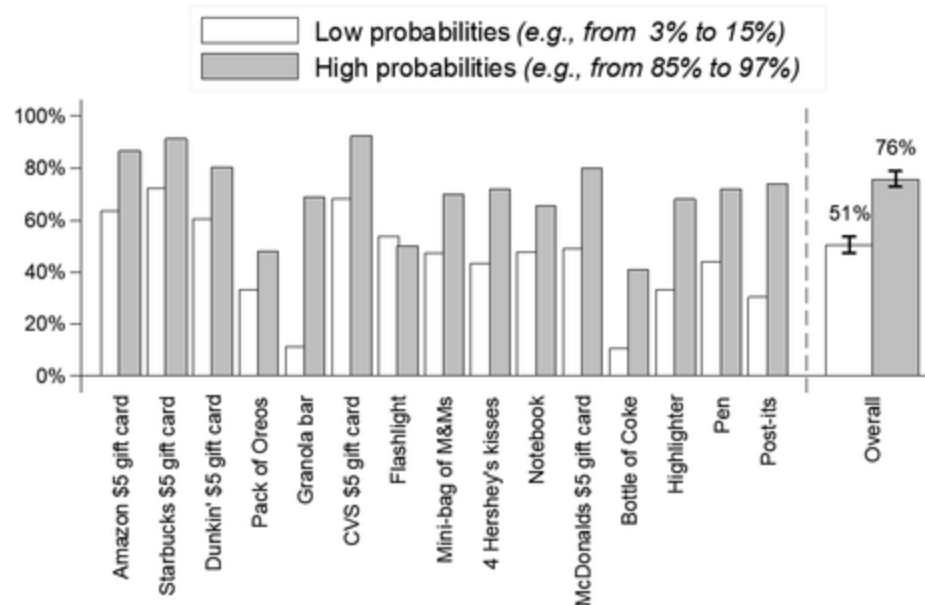


Figure 2. Percentage of decisions for which participants agreed to type “ab” the required amount of times in order to increase their probability of winning a prize, as a function of whether the win probability was high or low (manipulated within-subjects and across prizes). Bars to the left of the dotted line show the results for each prize (i.e., between-subjects results) and bars to the right show results collapsed across prizes (i.e., within-subjects results). Error bars to the right of the dotted line depict ± 1 clustered standard error.

Studies 3–5: Establishing the Generalizability of This Effect

We conducted three additional preregistered studies to establish the generalizability of the result found in Study 2 (that people are more likely to incur costs to increase high probabilities of success rather than low probabilities of success). In the service of making this article a reasonable length (or at least a less unreasonable length), we will give only a brief overview of the methods and results of each of these studies. The full write-ups of these studies are available in Online Supplements 3–5.

Study 3

In Study 2, we found that people were more likely to incur costs to increase high rather than low win probabilities. In Study 3, we examined whether our theory also applies to decisions to *maintain* win probabilities (at the expense of forgoing a possible bonus payment) rather than to *increase* win probabilities (at the expense of exerting effort). Prospective outcome bias predicts that people expect to be happier with a (costly) decision to maintain their initial chances of success when those initial chances of success are high than when they are low. Thus, people should be more inclined to refuse a benefit when it would come at the expense of reducing high rather than low chances of success.

In a design that was very similar to Study 2, laboratory participants in Study 3 made incentive-compatible decisions about whether to accept bonus payments in exchange for reduced chances of winning 15 prizes. For each prize (and, thus, within-subjects), the survey software randomly determined whether the potential win probabilities would be low (i.e., below 50%) or high (i.e., above 50%). For example, on one trial, participants were asked whether they would be willing to accept a \$0.50 bonus payment at the expense of reducing their probability of winning a \$5 Amazon gift card either from 16% to 7% (*low probabilities* condition) or from 93% to 84% (*high probabilities* condition).

After applying our preregistered exclusion rules, there were 103 participants and 728 observations in our main analysis. Consistent with prospective outcome bias, we found that participants were more likely to decline bonus payments that would have reduced high win probabilities than bonus payments that would have reduced low win probabilities, $b = .80$, clustered $SE = .28$, $p = .004$, $OR = 2.22$ (95% CI [1.29, 3.81]).

Study 4

In Studies 2 and 3, we found that people are more willing to increase/maintain high win probabilities than low win probabilities. In Study 4, we sought to further establish the generalizability of this effect, this time by more systematically sampling from the full range of probabilities (Wells & Windschitl, 1999). In addition to helping to establish the generalizability of the effect, this design allowed us to examine whether people's greater motivation to increase high win probabilities than low win probabilities varied according to the extremity of the potential probabilities and/or the size of the potential probability increase.

In the survey, online participants made hypothetical decisions about how much they would pay to increase their chances of winning each of 20 prizes (e.g., an iPhone X) from a baseline probability to an increased probability.⁸ For each prize (and, thus, within-subjects), the survey software randomly determined whether the potential win probabilities would be low (i.e., less than or equal to 50%) or high (i.e., greater than or equal to 50%). In the *low probabilities* condition, we sampled two probabilities from a uniform distribution between 0.01% and 50.00% (displayed to the nearest two decimal places). The lower of these two probabilities was assigned to be the baseline probability, and the higher of the two probabilities was assigned to be the increased probability. For the *high probabilities* condition, we determined the probabilities in the same way, but instead sampled from a uniform distribution between 50.00% and 99.99%.

After applying our preregistered exclusion rules, there were 392 participants and 7,781 observations in our main analysis. Replicating our effect from previous studies, we found that participants indicated a greater willingness to pay to increase high win probabilities than low win probabilities, $b = 18.13$, clustered $SE = 1.44$, $p < .001$.⁹ Additionally, we found that this effect was larger when the win probabilities were more extreme, $b = .81$, clustered $SE = .09$, $p < .001$. This result indicates that people are least inclined to pay to improve their chances of success when they are extremely unlikely to be happy with their decision (because success is extremely unlikely), and that they are most inclined to pay to improve their chances of success when they are extremely likely to be happy with their decision (because success is extremely likely). Thus, this result is consistent with prospective outcome bias. In general, when potential win probabilities are higher, prospective outcome bias predicts people will be willing to invest more; and that is what we found in Study 4.

Finally, we also found that the effect of higher win probabilities on willingness to pay was greater when the extent of the potential probability increase was also greater, $b = .32$, clustered $SE = .10$, $p = .002$. We speculate that this result arises because people may be unwilling to invest to attain an extremely small improvement in the expected value of a lottery regardless of their final win probability.

Study 5

Studies 2–4 demonstrated people’s greater preference for increasing high chances of success over low chances of success in incentive-compatible laboratory settings and

hypothetical online settings. In Study 5, we investigate the applicability of this finding to a variety of realistic, high stakes scenarios, including scenarios in which success constitutes obtaining positive outcomes and scenarios in which success constitutes avoiding negative outcomes.

For each of eight scenarios, online participants indicated whether they would take action to improve their chances of success. For each scenario (and, thus, within-subjects), the survey software randomly determined whether the potential probabilities of success would be favorable or unfavorable.

To test whether the preference for increasing favorable over unfavorable chances of success generalizes to a wide variety of potential successful outcomes, we constructed four scenarios that were about achieving success and four scenarios that were about avoiding failure. For example, one achieving success scenario read as follows, with the unfavorable probabilities version in the main text and the favorable probabilities version in brackets:

Imagine your business start-up is looking for investment, which it needs to be a sustainable source of income for you. There is a big, local investment fund that currently has a 5% [88%] chance of making an investment. You could add a new service to your start-up, which you know would appeal to the investment fund, but doesn't fit in your current business plan. If you add this new service, there is instead a 12% [95%] chance that the investment fund will invest in your start-up.

Would you add the new service?

Participants chose their answer from two options: “Yes: Add the new service to increase the chance of getting the investment to 12% [95%],” and “No: Keep the start-up as it is now and accept a 5% [88%] chance of getting the investment.”

One *avoiding failure* scenario read as follows, with the unfavorable probabilities version in the main text and the favorable probabilities version in brackets:

Imagine you are CEO of a small financial company, and your local government is considering new regulations that you think are deeply misguided, and could cost your company millions of dollars. Currently, there is around a 96% [19%] chance that these regulations are implemented. You could launch a lobbying campaign, which would reduce the chance to around 81% [4%]. However, this lobbying campaign would cost tens of thousands of dollars, and might come at a reputational cost.

Would you launch a lobbying campaign?

Participants chose their answer from two options: “Yes: Launch the campaign to reduce the chances that the misguided regulations are implemented to 81% [4%],” and “No: Forgo the campaign and accept a 96% [19%] chance that the misguided regulations are implemented.”

After applying our preregistered exclusion rules, there were 422 participants and 3,369 observations in our main analysis. Consistent with prospective outcome bias, we found that people were indeed more likely to incur costs to improve favorable probabilities than unfavorable probabilities, $b = .39$, clustered $SE = .07$, $p < .001$, $OR = 1.48$ (95% CI [1.29, 1.71]). We also preregistered to run the same regression separately for the *achieving*

success scenarios and the *avoiding failure* scenarios. The effect was significant for both the positive, *achieving success* outcomes, $b = .34$, clustered $SE = .10$, $p < .001$, $OR = 1.40$ (95% CI [1.16, 1.70]), and for the negative, *avoiding failure* outcomes, $b = .45$, clustered $SE = .10$, $p < .001$, $OR = 1.56$ (95% CI [1.29, 1.90]).

These results indicate that prospective outcome bias manifests for realistic, high-stakes scenarios. It also manifests whether people are striving to obtain positive outcomes or to avoid negative outcomes, and whether we frame the improvement in the chances of success as an increase in the probability of a positive outcome or a decrease in the probability of a negative outcome.

Study 6

Although the results thus far are wholly consistent with prospective outcome bias, the findings of our incentive-compatible studies are also consistent with what might be expected from anticipated regret (Zeelenberg, 1999) or waste aversion (Arkes, 1996). Research shows that people are more likely to anticipate feeling regret about a decision when they will get feedback on that decision, because the feedback may reveal that they would have been better off making a different decision (van de Ven & Zeelenberg, 2011; Zeelenberg & Pieters, 2004). In Study 2, the only way participants could learn that they would have been better off having made a different decision was if they had typed “ab” the required amount of times and had still failed to win a prize. Because this aversive situation was more likely to occur when the probability of winning was *low*, participants

might have been more willing to exert effort to improve their chances of winning when the probability of winning was *high*.

In Study 6, we adjusted Study 2's design to eliminate any differences in feedback across conditions. Specifically, we ensured that all participants knew for certain that they would find out whether their decisions influenced the results of the prize draws. This new design feature ensured that participants in the *low* versus *high probabilities* condition were equally likely to learn that their effort was inconsequential. Thus, if the results of the previous studies were driven purely by participants being averse to learning that their effort was wasted, then we would not expect those results to replicate in Study 6.

Method

Participants

Participants were paid \$10 for an hour-long laboratory session in a northeastern university, and this study was a 15-minute part of this session. We preregistered to collect participants from a week-long series of lab sessions, aiming to get 150 participants in total. We also preregistered to include only the original response of any participant who completed the study more than once, resulting in two exclusions. After all preregistered exclusions, including some described below, our final sample comprised 181 participants (mean age = 26.4, 61.9% female).

Design

As in Study 2, participants decided whether to improve their chances of winning each of 15 prizes by typing “ab” on a keyboard a given amount of times. For each prize (and so, within subjects), the survey software randomly determined whether the potential win probabilities would be low (i.e., below 50%) or high (i.e., above 50%). However, in Study 6, participants knew at the time of making their decisions that they would find out whether those decisions influenced the outcome of the prize draws.

Procedure

At the beginning of the survey, we explained the task to participants just as we had done in Study 2. However, at the end of these instructions, we added one additional instruction, which explained to participants that they would find out whether their decisions influenced the outcome of the prize draw:

On this page, you will learn more details about how the prize draw works. Specifically, the computer draws a random number from 1 to 100. If this number is less than or equal to your probability of winning, then you win the prize. At the end of the survey, we will tell you what random number the computer generated, and we will inform you of whether your decision about whether to type “ab” influenced whether or not you won the prize. Thus, if you choose to type “ab”, you will learn whether or not you won the prize because of this decision. If you choose not to type “ab”, you will learn whether or not you failed to win the prize because of this decision.

Then, participants answered the same two comprehension check questions as in Study 2, along with a third, additional comprehension check in which they had to confirm

that they understood that they would find out whether their decision influenced the outcome of the prize draw.

The procedure from this point onward was identical to Study 2, except for two changes. First, at the top of each of the screens on which participants made their decisions, we reminded them that they would find out whether their decisions influenced the outcome. Specifically, we told them:

Please answer the question below. Remember, whatever you choose, you will find out whether and how your answer to this question influences the result of the prize draw.

The second change was that when participants learned the results of the prize draw, we told them whether their decision had influenced the outcome.¹⁰

Results and Discussion

We preregistered two additional exclusion rules to ensure that the participants were motivated to win the prizes and understood the instructions sufficiently. First, we excluded all decisions from any participant who failed to pass all three comprehension check questions within two attempts. Second, as in Study 2, we also excluded decisions about prizes for which participants rated their motivation to win at less than a 4 (*neutral*) on a 7-point scale (1 = *extremely unmotivated* to 7 = *extremely motivated*). This left us with 181 participants and 1,526 observations for our main analysis.

As in the previous studies, we examined whether participants were more likely to invest effort to improve already favorable chances than unfavorable chances. We

preregistered to test this hypothesis using a binary dependent variable to indicate whether the participant decided to type “ab” the required amount of times to increase the win probability (1 = yes, 0 = no). Thus, each participant provided up to 15 observations of the dependent variable, one for each prize for which they rated their motivation as at least “neutral”. Using a logit model, we regressed the dependent variable on a *high probabilities* condition dummy variable (1 = *high probabilities* condition; 0 = *low probabilities* condition), and fixed effects for each prize, clustering standard errors by participant.

Figure 3 displays the results. We found that people were indeed more likely to invest effort to improve high win probabilities than low win probabilities, $b = .74$, clustered $SE = .17$, $p < .001$, $OR = 2.11$ (95% CI [1.52, 2.91]). This finding cannot be explained by regret aversion or by a general aversion to wasted effort, because participants knew that the probability of finding out that their decision did not influence the outcome was the same in the *high probabilities* condition as in the *low probabilities* condition. However, this finding does follow from prospective outcome bias. At high probabilities of winning, participants know that their investment of effort is likely to seem justified (by a positive outcome), so they are motivated to improve their chances. At low probabilities of winning, participants know that their investment of effort is *unlikely* to seem justified (by a positive outcome), so they are less motivated to improve their chances.

Study 6: Participants were still more likely to expend effort to increase high win probabilities than low win probabilities when they knew they would find out whether their effort was pivotal




Figure 3. Percentage of decisions for which participants agreed to type “ab” the required amount of times in order to increase the probability of winning a prize, as a function of whether the win probability was high or low (manipulated within-subjects and across prizes). Bars to the left of the dotted line show results for each prize (i.e., between-subjects results) and bars to the right show results collapsed across prizes (i.e., within-subjects results). Error bars to the right of the dotted line depict ± 1 clustered standard error.

Study 7

In Studies 2–6, we found that people were more likely to work to increase or maintain high probabilities of success than low probabilities of success. In all of these studies, we constructed the stimuli so as to ensure that the low probabilities and the high probabilities were equally extreme (i.e., equally far from 50%). In so doing, we effectively ruled out that this finding could be driven by a general overweighting of very small probabilities that is perfectly mirrored by a general underweighting of very large probabilities.

However, some scholars have suggested that people may weigh low probabilities differently from how they weigh high probabilities (Delqu   & Cillo, 2006; Kahneman & Tversky, 1979). If this is true, then it is possible to construct a scenario in which probability weighting could account for our findings. Specifically, if the probability weighting function is *steeper* for high probabilities than for low probabilities—that is, if people are *more sensitive* to differences in high probabilities than to differences in low probabilities—then we would expect people to be more motivated to *increase* high win probabilities than low win probabilities.

In Study 7, we pitted this *asymmetric probability weighting* explanation of our results against our prospective outcome bias explanation. To accomplish this, we again asked participants to indicate how much they would be willing to pay to increase high probabilities of winning a prize draw (e.g., from 70% to 80%) or low probabilities of winning a prize draw (e.g., from 20% to 30%), but we also asked them to *separately* indicate their willingness to pay to enter prize draws at the two low win probabilities (e.g., 20% and 30%) and at the two high win probabilities (e.g., 70% and 80%). If our findings result from people being more sensitive to the difference between high win probabilities than low win probabilities, then all three of the following results should emerge:

1. As observed in the previous studies, participants should pay more to increase high win probabilities than low win probabilities.

2. The difference between how much participants are *separately* willing to pay to enter the two high-probability prize draws should be greater than the difference between how much participants are *separately* willing to pay to enter the two low-probability prize draws.

3. The size of Result 2 should equal the size of Result 1.

It is important to note that these same predictions also follow from another alternative explanation for our findings. Some research on the goal gradient effect suggests that goals serve as reference points that operate according to the principles of Prospect Theory (Heath, Larrick, & Wu, 1999). Because Prospect Theory holds that people are diminishingly sensitive to increasing losses, this account predicts that people will be more motivated to make progress toward a goal as they get closer to the goal. If one assumes that people adopt a goal of a 100% win probability, then this account predicts that people are more sensitive to differences between higher probabilities (closer to the goal) than lower probabilities (further from the goal). This account thus makes identical predictions to those of the asymmetric probability weighting account. Specifically, as well as predicting greater willingness to improve higher win probabilities (like prospective outcome bias), it also predicts that greater sensitivity to higher probabilities should show up just as forcefully in differences between people's willingness to pay to enter separate prize draws (*unlike* prospective outcome bias).¹¹

In contrast, prospective outcome bias predicts that people will pay more to increase high win probabilities than to increase low win probabilities, but it does not predict that

this effect will be as strong when comparing (a) the difference between what people are *separately* willing to pay for the two high-probability prize draws and (b) the difference between what they are *separately* willing to pay for the two low-probability prize draws. To understand why, recall that prospective outcome bias asserts that the value that people place on an *improvement* in a lottery is driven at least in part by the *expected outcome* of that lottery. This is because the outcome of the lottery can be used to justify the costs incurred to acquire the improvement. Consequently, although people's valuation of an improvement from an 80% to a 90% win probability is partially driven by the appeal of a 10% increase in their chances of winning, it will also be partially driven by the (relatively high) appeal of a 90% win probability. Similarly, although people's valuation of an improvement from a 10% to a 20% win probability is partially driven by the appeal of a 10% increase in their chances of winning, it will also be partially driven by the (relatively low) appeal of a 20% win probability. Thus, when faced with a decision about whether to incur a cost so as to *improve* chances of success, people will be more likely to do so when those improved chances would be 90% than when they would be 20%. This is what we have found thus far.

Importantly, however, prospective outcome bias predicts no such effect when people are deciding whether to enter separate lotteries. To understand why, first consider that deciding whether to *enter* a lottery is equivalent to deciding whether to *improve* one's win probability from 0%. Thus, the expected value of the improvement in win probability from entering the lottery is identical to the expected value of the lottery itself. Consequently, prospective outcome bias predicts that a person's valuation of that lottery

will be driven solely by its expected value. For example, it predicts that a risk-neutral person will pay \$1 more for a 90% chance of winning \$10 (expected value = \$9) than for an 80% chance of winning \$10 (expected value = \$8), but will also pay \$1 more for a 20% chance of winning \$10 (expected value = \$2) than for a 10% chance of winning \$10 (expected value = \$1).¹² In other words, although prospective outcome bias predicts that people will be biased toward *improving* higher probabilities of success, it does not predict that there will generally be greater differences between people's separate valuations of two high probabilities of success than between their separate valuations of two low probabilities of success.¹³

Method

Participants

Participants were paid \$10 for an hour-long laboratory session in a northeastern university, and this study was a 15-minute part of this session. We preregistered to collect participants from a week-long series of lab sessions, aiming to get 200 participants in total. We also preregistered to include only the original response of any participant who completed the study more than once, resulting in three exclusions. After all preregistered exclusions, including some described below, our final sample comprised 140 participants (mean age = 21.7, 65.0% female).

Design

Participants decided how much they would be willing to pay to increase their probabilities of winning each of 10 prize draws (probability increase questions), and to

enter each of 20 prize draws (separate probabilities questions). Half of the prize draws in each condition had low win probabilities (i.e., below 50%) and the other half had high win probabilities (i.e., above 50%). There were five different prizes, and for each prize, participants made six decisions. Specifically, for each prize, participants indicated their willingness to pay to (decision 1) increase a low win probability (e.g., from 3% to 9%), (decision 2) increase a high win probability (e.g., from 91% to 97%), and (decisions 3–6) enter a prize draw at each of the four possible win probabilities (e.g., 3%, 9%, 91%, and 97%). For example, if participants could increase their win probabilities from 3% to 9% or from 91% to 97% for a given prize, the same participants would also indicate their willingness to pay to enter draws for the same prize at each of 3%, 9%, 91% and 97%. For each participant's decisions about each prize, we randomly drew one set of probabilities from four possible sets: {{ '{' }}3%, 9%, 91%, 97%}; {{ '{' }}6%, 12%, 88%, 94%}; {{ '{' }}6%, 15%, 85%, 94%}; and {{ '{' }}11%, 20%, 80%, 89%}.

Procedure

At the beginning of the survey, we explained to participants that they would make 30 decisions about how much they would be willing to pay either to increase the probabilities of winning prize draws that they were automatically entered into, or to enter prize draws at given win probabilities. To ensure that participants took their decisions seriously, we truthfully told them that we would randomly select one of these 30 prize draws to conduct for real, and we displayed a selection of the prizes on a table at the front of the room. We determined whether or at what probability participants were entered into

the selected prize draw using a procedure that was designed to incentivize participants to accurately report their willingness to pay (Becker, Degroot, & Marschak, 1964).¹⁴

After explaining this procedure to participants, we encouraged them to give granular responses by asking them to respond to all of the questions in cents rather than dollars. On this page, we asked them to write \$9.87 in cents as an example and allowed them to proceed only after they had successfully entered “987”. Participants then answered a comprehension question about the instructions, in which they had to confirm that the survey would randomly select one of the prize draws to conduct for real. On the next page they rated their motivation to win each of the five prizes.

At this juncture, the survey software counterbalanced whether the participants first answered the 10 questions about increasing win probabilities or the 20 questions about entering prize draws at given probabilities. Before starting the block of questions about increasing win probabilities, participants answered a comprehension check question in which they had to confirm that they were indeed answering questions about *increasing* their chances of winning a prize, as opposed to *entering* into a prize draw at a given probability. Participants then reported their willingness to pay to increase their win probabilities, with each question presented one at a time on the computer screen and in a random order. For each question, we informed participants of the prize, what their baseline probability of winning would be, and what their increased probability of winning would be if they reported a sufficiently high willingness to pay. See Figure 4 for an example of what participants saw on an Amazon gift card trial.

Screenshot of separate probabilities question (corresponding to baseline probability of winning)

What is the most that you would pay in order to have a **88%** chance of winning a **\$5 Amazon gift card**?

Please answer in **CENTS**, not dollars.

Screenshot of separate probabilities question (corresponding to increased probability of winning)

What is the most that you would pay in order to have a **94%** chance of winning a **\$5 Amazon gift card**?

Please answer in **CENTS**, not dollars.

Screenshot of probability increase question

Imagine you were already in the following prize draw.

Prize: \$5 Amazon gift card

Probability of winning: 88%

What is the most that you would pay in order to increase your probability of winning from **88% to 94%**?

Please answer in **CENTS**, not dollars.

Figure 4. Screenshots of the *high probabilities* condition for the Amazon gift card trial in Study 7.

Similarly, before starting the block of questions about *entering* prize draws, participants answered a comprehension check question in which they had to confirm that they were indeed answering questions about *entering* into a prize draw at a given probability, as opposed to *increasing* their chances of winning a prize. Participants then reported their willingness to pay to enter the prize draws, with each question presented one at a time on the computer screen and in a random order. For each question, we informed

participants both of the prize and of what their probability of winning would be if they were entered into the prize draw (see Figure 4).

After participants made their decisions, the survey randomly selected one of their responses to implement for real, informed participants both of which response had been selected and of the implications of this response for their chances of winning the relevant prize, and finally, revealed the outcome of the prize draw (if they were entered into it). Participants were paid as in Studies 2, 3, and 6.

Results and Discussion

We preregistered two additional exclusion rules to ensure that participants were motivated to win the prizes and understood the instructions sufficiently. First, we excluded all decisions from any participant who failed to pass all three comprehension check questions within two attempts. Second, as in Studies 2, 3, and 6, we also excluded decisions about prizes for which participants rated their motivation to win at less than a 4 (*neutral*) on a 7-point scale (1 = *extremely unmotivated* to 7 = *extremely motivated*). This left us with 140 participants and 1,724 observations for our main analysis.

For each item/participant, we analyzed four key variables: (a) their stated willingness to pay to increase high win probabilities, (b) their stated willingness to pay to increase low win probabilities, (c) the difference between how much they were separately willing to pay to enter the prize draw at the two high win probabilities, and (d) the difference between how much they were separately willing to pay to enter the prize draw at the two low win probabilities. Variables (a) and (b) represent

participants' *stated* willingness to pay to increase high and low win probabilities, whereas variables (c) and (d) represent their *implied* willingness to pay to increase high and low win probabilities.

If people are more motivated to increase high than low probabilities because they are generally more sensitive to high than low probabilities, then the difference between participants' stated willingness to pays (i.e., (a) and (b)) should be the same as the difference between their implied willingness to pays (i.e., (c) and (d)). If this effect is instead driven by prospective outcome bias, then the difference between participants' stated willingness to pays (i.e., (a) and (b)) should be greater than the difference between their implied willingness to pays (i.e., (c) and (d)).

Each participant provided up to 20 observations: for each of the five prizes, they contributed four observations, one for each cell of the 2 (separate probabilities vs. probability increase) \times 2 (low probabilities vs. high probabilities) design. To limit the effects of outliers, we preregistered to winsorize this dependent variable at the 5th and 95th percentiles of all observations; that is, all observations below the 5th percentile were treated as equal to the 5th percentile, and all observations above the 95th percentile were treated as equal to the 95th percentile. We regressed this dependent variable on a *high probabilities* condition contrast-coded variable ($-0.5 = \text{low probabilities}$ condition; $0.5 = \text{high probabilities}$ condition), a *probability-increase* condition contrast-coded variable ($-0.5 = \text{separate probabilities}$ condition; $0.5 = \text{probability increase}$ condition), and their interaction, clustering standard errors by participant and including fixed effects for each item.

Figure 5 displays the results. Replicating the previous studies, there was a highly significant effect of the *high probabilities* condition, $b = 25.79$, clustered $SE = 3.09$, $p < .001$, indicating that participants paid more to increase high win probabilities than to increase low win probabilities. There was also a highly significant effect of the *probability increase* condition, $b = 40.23$, clustered $SE = 4.48$, $p < .001$, indicating that participants' *stated* willingness to pay to increase their win probabilities exceeded their *implied* willingness to pay to increase those win probabilities. Most important, however, was that the High Probabilities \times Probability Increase interaction was positive and highly significant, $b = 39.03$, clustered $SE = 5.18$, $p < .001$.¹⁵ This interaction indicates that the effect of greater willingness to pay to increase high versus low win probabilities was much weaker when we measured willingness to pay for the probability increase indirectly; that is, by differencing participants' separate reports of their willingness to pay to enter prize draws at the potential improved win probability and at the initial unimproved win probability. Indeed, although participants in Study 7 paid significantly more to increase high rather than low win probabilities when they were asked directly, $b = 45.30$, clustered $SE = 4.44$, $p < .001$, they did not exhibit such a strong (or even significant) tendency in the difference between their separate reports of how much they would be willing to pay for the improved and unimproved win probabilities, $b = 6.27$, clustered $SE = 3.59$, $p = .083$.

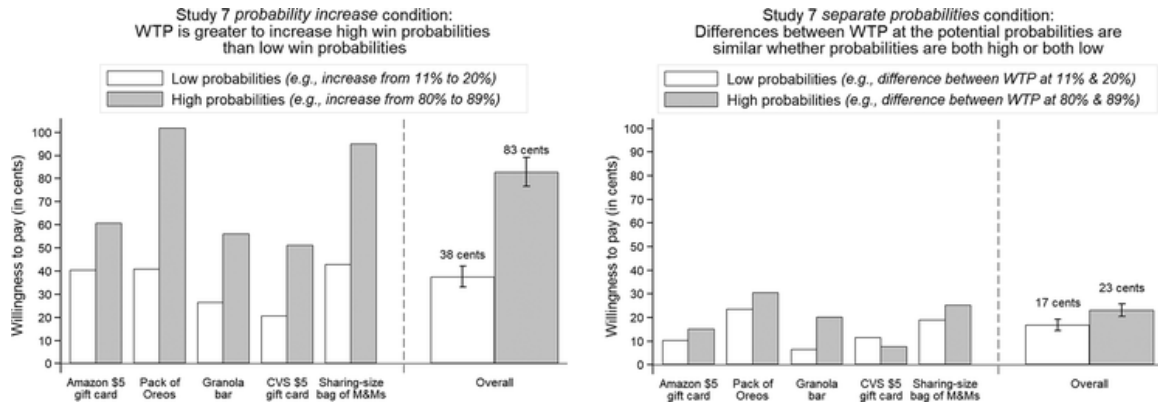


Figure 5. The left panel exhibits participants' willingness to pay (WTP) to increase a win probability, as a function of whether the potential probabilities were both high or both low (manipulated within-subjects). The right panel exhibits the difference between participants' WTP for a higher probability of winning a prize and their WTP for a lower probability of winning a prize, as a function of whether the potential probabilities were both high or both low (manipulated within-subjects). Within each panel, bars to the left of the dotted line show the results for each prize and bars to the right of the dotted line show results collapsed across prizes. Error bars to the right of the dotted line depict ± 1 clustered standard error.

In sum, if people are more sensitive to differences between high probabilities than to differences between low probabilities, this effect is small, and cannot explain the very strong tendency to be more motivated to increase high versus low win probabilities that we have observed in our studies. Thus, the results of Study 7 are consistent with prospective outcome bias, but not an alternative explanation based on participants' asymmetric sensitivity to high versus low probabilities (Delquie & Cillo, 2006; Kahneman & Tversky, 1979).

Study 8

We designed Study 8 to more directly test whether prospective outcome bias can explain the effects in Studies 2–7, in which people were more motivated to increase high rather than low chances of success. According to prospective outcome bias, this effect

arises because people anticipate how outcome bias will make them feel about a costly decision: that is, a successful outcome will make the decision seem justified and thus the costs of such a decision feel *less* painful, but an unsuccessful outcome will make the decision seem misguided and thus make the costs of the decision feel *more* painful. Importantly, people know that they are more likely to feel better about incurring costs to improve high rather than low chances of success, and thus, they are more likely to do so.

If this account is correct, then individual differences in the extent of outcome bias should predict the preference for improving more favorable chances of success. Specifically, people who exhibit greater outcome bias with respect to a decision to improve their chances of success should also exhibit a greater preference for increasing high chances of success over low chances of success. We test this prediction in Study 8 using hypothetical scenarios that were very similar to those that we used in Study 5. For each scenario, we asked participants two sets of questions. The first set measured their relative preference for increasing favorable over unfavorable probabilities of success, and the second set measured outcome bias with respect to the decision to improve chances of success. If prospective outcome bias is driving the results, then participants who show the most outcome bias for a specific scenario should also be most likely to show greater motivation to improve high versus low chances of success for that scenario.

Method

Participants

We recruited participants from Amazon’s Mechanical Turk (MTurk). Participants received \$1.00 for completing the study. We preregistered to collect data from 400 participants. In the event of multiple responses from a single MTurk ID or IP address, we preregistered to include the original response only, resulting in 64 exclusions, and a final sample of 373 participants (mean age = 38.2, 46.0% female).

Design

The scenarios we used in Study 8 were the same as those we used in Study 5, with slight modifications (see Table 2 for a summary of these scenarios). However, to keep the survey at a reasonable length, participants were randomly presented with only six of the eight scenarios. For each scenario, participants answered two sets of questions. First, they answered two questions designed to measure their relative preference for increasing favorable over unfavorable chances of success.¹⁶ Second, they answered two questions designed to measure outcome bias with respect to the decision to improve chances of success.¹⁷ Specifically, we asked participants how they would feel about the decision *to take action* to improve their chances of success, both given that the *preferred* outcome occurred and given that the *less* preferred outcome occurred.

Table 2
Scenarios Used in Studies 5 and 8

Focal outcome	Required action to improve probabilities	Possible change in probability ^a	
		Unfavorable probabilities	Favorable probabilities
Achieving success			
Successfully suing	Hiring an expensive lawyer	From 5% to 8%	From 92% to 95%
Winning local state legislature	Donation to local party	From 4% to 11%	From 89% to 96%
Investment for a start-up	Adding a new service	From 5% to 12%	From 88% to 95%
Publishing a book	Hiring an editor	From 5% to 8%	From 92% to 95%
Avoiding failure			
Patient fatality	Painful treatment	From 96% to 91%	From 9% to 4%
Product imitation by rival	File a patent	From 97% to 93%	From 7% to 3%
Laptop breaks	Expensive PC repair store	From 95% to 78%	From 22% to 5%
Misguided regulations	Lobbying campaign	From 96% to 81%	From 19% to 4%

^a Probabilities are expressed in this table in terms of the focal outcome, as they were displayed to participants. So, for positive, *achieving-success* outcomes, the probabilities refer to the likelihood of the positive outcome, and for negative, *avoiding-failure* outcomes, the probabilities refer to the likelihood of the negative outcome.

Procedure

At the beginning of the survey, before participants read any scenarios, we told them that they would judge how likely they would be to take action to improve the chances of various outcomes. On the next page, participants answered an attention check question in which they had to confirm this instruction. If they failed to answer correctly on their first attempt, the survey prevented them from continuing and we collected no data from them.

Participants then saw six of the eight possible scenarios, and after each scenario, answered the two sets of questions. For example, some participants read the following scenario:

Imagine your business start-up is looking for investment, which it needs to be a sustainable source of income for you. There is a big, local investment fund that is considering making an investment in your start-up. You could add a new service to your

start-up to increase the chance of getting the investment. You know the new service would appeal to the investment fund, but it does not fit in your current business plan.

On the same page as this scenario, we then asked the first set of questions, designed to measure the relative preference for increasing favorable over unfavorable chances of success. Specifically, we asked participants the following questions at both unfavorable probabilities (used in the main text) and favorable probabilities (in brackets), with order counterbalanced between subjects. Participants answered on a 7-point scale from 1 = *very unlikely* to 7 = *very likely*:

Imagine that there is a 5% [88%] chance that the investment fund will invest in your start-up without the new service, and a 12% [95%] chance that the investment fund will invest in your start-up if you add the new service. How likely would you be to add the new service?"

Over the next two pages, we then asked the second set of questions, designed to measure participants' susceptibility to outcome bias for the decision to act to improve their chances of success. First, we told participants to imagine that they added the new service, and then on a 7-point scale from 1 = *very bad* to 7 = *very good*, we asked them how they would feel about this decision given that the desired investment was attained, and then how they would feel about this decision given that the desired investment was *not* attained.¹⁸

After answering these two sets of questions for the six scenarios, participants entered their demographic information.

Results and Discussion

Following our preregistration, we measured the relative preference for increasing high over low chances of success using the first set of questions for each scenario. Specifically, we calculated the difference between participants' rating of the likelihood that they would take action at *favorable* probabilities of success and the equivalent rating at *unfavorable* probabilities of success. We regressed this dependent variable on a constant term, and we clustered standard errors. Consistent with our previous studies, we found that participants preferred to improve favorable probabilities over unfavorable probabilities, $b = .44$, clustered $SE = .05$, $p < .001$.

Importantly, we predicted that this preference would be greater for participants who exhibited greater outcome bias with respect to the decision to improve chances of success. We measured outcome bias by calculating the difference between the participants' rating of how they would feel if they acted and achieved a successful outcome from the equivalent rating for if they acted and did *not* achieve a successful outcome. We regressed our dependent variable on this measure, including fixed effects for scenario and clustering standard errors by participant.

As predicted, we found that participants who were more prone to outcome bias for a particular scenario were also more likely to be motivated to increase favorable versus unfavorable probabilities, $b = .17$, clustered $SE = .02$, $p < .001$. Figure 6 shows that the tendency to be more motivated to increase favorable versus unfavorable probabilities was observed most strongly for those with the greatest outcome bias, and that this tendency was absent among those who exhibited no outcome bias. This pattern supports our theory of prospective outcome bias. When people expect to judge a decision more favorably after

success than after failure, they are also more motivated to increase high rather than low probabilities of success. And, when they do not expect to judge a decision more favorably after success than after failure, they are *not* more motivated to increase high rather than low probabilities of success.

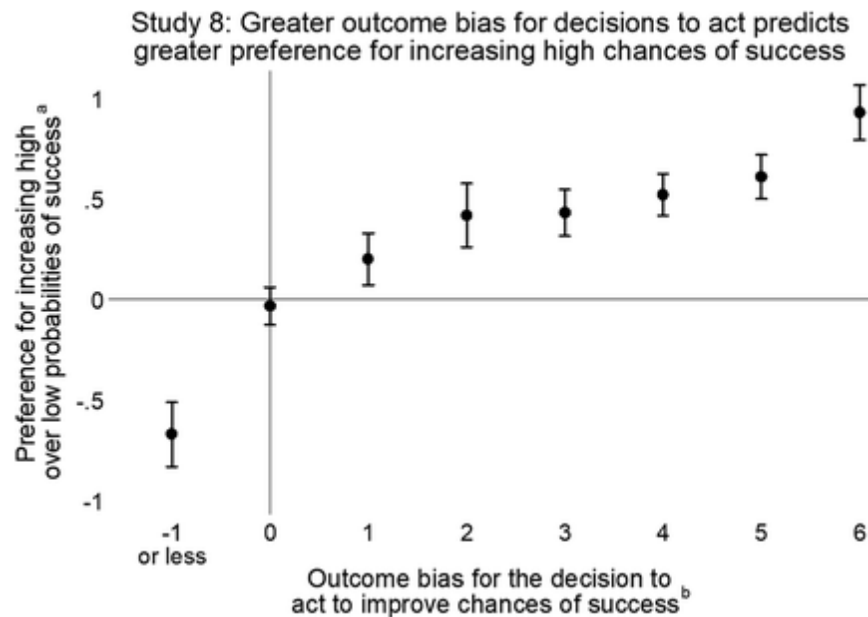


Figure 6. Participants' relative preference for increasing favorable chances of success in a given scenario over unfavorable chances of success, plotted as a function of their outcome bias with respect to the decision to increase chances of success in that scenario. Each participant provided 6 observations, one for each of a randomly determined 6 of the 8 scenarios in Table 2. Error bars depict ± 1 clustered standard error. ^aThe preference for increasing favorable over unfavorable chances of success is measured by the difference between participants' ratings of how likely they would be to improve their chances of success (on a scale from 1 = "Very unlikely" to 7 = "Very likely") at favorable chances versus at unfavorable chances. ^bOutcome bias with respect to the decision to act is measured by the difference between participants' rating of how they would feel about a decision to act (on a scale of 1 = "Very bad" to 7 = "Very good") after achieving a successful outcome versus an unsuccessful outcome.

Study 9

Studies 2–8 establish that people are more willing to increase chances of success when those chances are already high, and that prospective outcome bias is the likely explanation of this effect. In Study 9, we examine a different prediction of prospective outcome bias: people will prefer to increase the probability of winning a larger reward by a smaller amount than to increase the probability of winning a smaller reward by a larger amount.

This prediction follows from the notion that potentially unnecessary costs will feel more justified, and thus less painful, if the decision to incur those costs is accompanied by a more positive outcome (i.e., a larger prize). To illustrate, imagine that someone invests to increase their chances of winning \$100 from 89% to 90% (and so the increase in the expected value of the lottery is $\$1 = (90\% - 89\%) \times \100). This person is likely to win a large sum of \$100, and so they are also likely to feel extremely good about the outcome and thus the decision to invest. Consequently, when given this opportunity, people who anticipate their evaluations of their decisions will be very inclined to invest. Now instead imagine that someone invests to increase their chances of winning a smaller sum of \$10 by a greater amount, from 80% to 90% (and so the increase in the expected value of the lottery is again $\$1 = (90\% - 80\%) \times \10). This person is likely to win a small sum of \$10, and so they are also likely to feel good-but-not-great about the outcome and thus the decision to invest. Consequently, when given this opportunity, people who anticipate their evaluations of their decisions will be only moderately inclined to invest. Because people's decision evaluations, and thus their *anticipated* decision evaluations, are tied to outcomes rather than to their decisions' effect on the probability of success, a person who invests to increase

the probability of winning \$100 by 1% is likely to anticipate feeling better about his or her decision than a person who invests to increase the probability of winning \$10 by 10%. Thus, holding constant the potential increase in the expected value of the lottery, prospective outcome bias predicts that people will be more likely to invest to acquire a *smaller* improvement in the chances of winning a *larger* prize than to acquire a *larger* improvement in the chances of winning a *smaller* prize.¹⁹

Method

Participants

Participants were paid \$10 for an hour-long laboratory session in a northeastern university, and this study was a 15-minute part of this session. We preregistered to collect participants from a week-long series of lab sessions, aiming to get 150 participants in total. We also preregistered to include only the original response of any participant who completed the study more than once, resulting in two exclusions. After all preregistered exclusions, including some described below, our final sample comprised 114 participants (mean age = 23.5, 56.1% female).

Design

Participants made 16 decisions about whether to improve their chances of winning various monetary prizes by typing “ab” on a keyboard a given amount of times. For eight of these decisions, the potential win probabilities were low (i.e., below 50%) and for the other eight of these decisions, the potential win probabilities were high (i.e., above 50%). Within both the *high* and the *low probabilities* conditions, four of the decisions involved

a *large* potential win probability increase of 25% (for relatively *small* monetary prizes) and the other four decisions involved a *small* potential win probability increase of 5% (for relatively *large* monetary prizes).

To ensure that the increase in the expected value of the prize draw from investing effort was the same across conditions, each monetary prize in the 25% *probability increase* condition corresponded to a monetary prize in the 5% *probability increase* condition that was five times larger. For example, the first row of Table 3 shows that, on one trial in the 25% *probability increase* condition, participants were asked whether they would be willing to type “ab” 50 times to increase their probability of winning \$0.10 from 4% to 29% (*low probabilities* condition) or from 71% to 96% (*high probabilities* condition). Thus, the increase in the expected value of the prize draw from typing “ab” 50 times was always $\$0.025 = (29\% - 4\%) \times \0.10 . The fifth row of Table 3 shows that, on the corresponding trials in the 5% *probability increase* condition, participants were asked whether they would be willing to type “ab” 50 times to increase their probability of winning \$0.50 ($= 5 \times \0.10) by a smaller amount such as from 14% to 19% (*low probabilities* condition) or from 81% to 86% (*high probabilities* condition). Thus, as for the decisions detailed in line 1 of Table 3, the increase in the expected value of the prize draw from typing “ab” 50 times was also $\$0.025 = (19\% - 14\%) \times \0.50 . To control for probability weighting, we sampled the 5% probability increases to span the same range of probabilities as the 25% probability increases. For example, if a participant was assigned a potential 25% probability increase from 4% to 29% for one of the prizes,

the corresponding 5% probability increase for that prize would be sampled from one of: 4%–9%, 9%–14%, 14%–19%, 19%–24%, and 24%–29%.

Table 3
Study 9 Stimuli

Prize ^a	Required instances of typing “ab” to increase win probability	Possible increase in win probability	
		Low probabilities	High probabilities
25% probability increase (small prizes)			
\$.10	50	From 4% to 29%	From 71% to 96%
\$.20	75	From 5% to 30%	From 70% to 95%
\$.25	100	From 6% to 31%	From 69% to 94%
\$.40	125	From 7% to 32%	From 68% to 93%
Example of randomly sampled 5% increase in win probability ^b			
5% probability increase (large prizes)			
\$.50	50	From 14% to 19%	From 81% to 86%
\$1.00	75	From 15% to 20%	From 80% to 85%
\$1.25	100	From 16% to 21%	From 79% to 84%
\$2.00	125	From 17% to 22%	From 78% to 83%

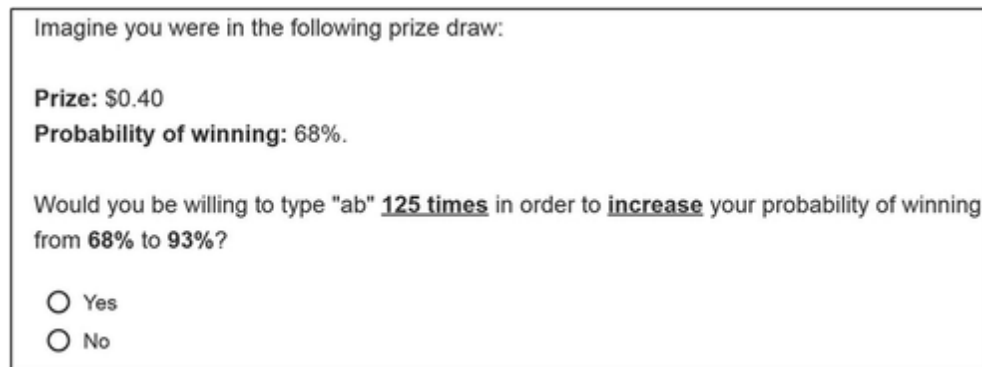
^a The first, second, third, and fourth prizes in the 5% probability increase condition are each five times the value of the first, second, third, and fourth prizes in the 25% probability increase condition, respectively. Thus, the potential increase in the expected value of the prize draw was controlled across conditions. ^b For each decision in the 5% probability increase condition, we randomly sampled the 5% probability increase from a set of five probability increases spanning the full range of the corresponding 25% probability increase (to control for probability weighting). For example, for the decision in the 5% probability increase condition (fifth row) corresponding to the 25% probability increase from 4% to 29% (first row), we randomly sampled from the following probability increases: 4% to 9%, 9% to 14%, 14% to 19%, 19% to 24%, and 24% to 29%.

Procedure

At the beginning of the survey, we explained to participants that they would make 16 decisions about whether to improve their chances of winning various monetary prizes by typing “ab” on a keyboard a given amount of times. To ensure that participants took their decisions seriously, we truthfully told them that we would randomly select one of these 16 prize draws to conduct for real. Participants then answered two comprehension questions about these instructions, paralleling those of Study 2.

Immediately after responding to the comprehension questions, participants practiced typing “ab” 100 times. They then made their decisions, which were presented one at a time on the computer screen and in a random order. For each decision, we informed participants of the monetary prize, what their baseline probability of winning would be, and what their increased probability of winning would be if they agreed to type “ab” the required amount of times. See Figure 7 for an example of what participants saw on the Amazon gift card trial and see Table 3 for details of each prize draw.

Participants were paid as in Study 2. At the end of the survey, participants entered their demographic information.



Imagine you were in the following prize draw:

Prize: \$0.40
Probability of winning: 68%.

Would you be willing to type "ab" **125 times** in order to **increase** your probability of winning from **68%** to **93%**?

☐ Yes
☐ No

Figure 7. Screenshot of \$0.40 trial for the *high probabilities/25% probability increase* conditions in Study 9.

Results and Discussion

We preregistered an additional exclusion rule to ensure that the participants understood the instructions. Specifically, we excluded all data from any participant who failed the first comprehension question on their first attempt, leaving us with 114 participants and 1,824 observations for our main analysis.

We had two hypotheses. First, as in the previous studies, we predicted that people are more likely to invest effort to improve already favorable chances than unfavorable chances. Second, we predicted that people are more likely to invest effort to incrementally improve their chances of high stakes outcomes than to substantially improve their chances of low stakes outcomes. To test our hypotheses, we preregistered to use a binary dependent variable to indicate whether the participant decided to type “ab” the required amount of times to increase the probability of winning the bonus payment (1 = *yes*; 0 = *no*). Using a logit model, we regressed this dependent variable on a *high probabilities* condition contrast-coded variable (-0.5 = *low probabilities* condition; 0.5 = *high probabilities* condition), a *5% probability increase* contrast-coded variable (-0.5 = *25% probability increase* condition; 0.5 = *5% probability increase* condition), and their interaction. In addition, we included fixed effects for each of the four possible increases in expected value of the prize draw that would have been accrued from typing “ab” the required amount of times, and we clustered standard errors by participant.

As shown in Figures 8 and 9, both of our hypotheses were supported. First, as in the previous studies, we found that participants were more likely to invest effort to improve high win probabilities than low win probabilities, $b = .56$, clustered $SE = .08$, $p < .001$, $OR = 1.75$ (95% CI [1.48, 2.07]; see Figure 8). Second, and also consistent with prospective outcome bias, we found that participants were indeed more likely to increase the probability of winning large bonuses by 5 percentage points than to increase the probability of winning smaller bonuses by 25 percentage points, $b = .86$, clustered $SE = .13$, $p < .001$, $OR = 2.37$ (95% CI [1.84, 3.07]; see Figure 9). This result corroborates the

second key prediction of prospective outcome bias: when deciding whether an improvement in the chances of success is worth the cost, people put more weight on the relative desirability of a successful outcome than they do on how much their chances of attaining it would be increased.

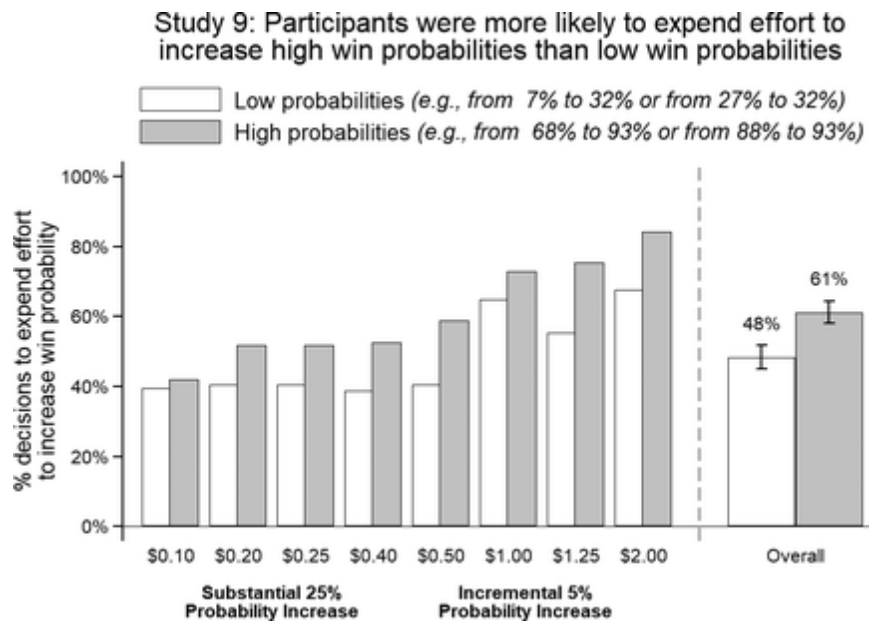


Figure 8. Percentage of decisions for which participants agreed to type “ab” the required amount of times in order to increase the probability of winning a monetary prize, as a function of whether the win probability was high or low (within-subjects for each prize). Bars to the left of the dotted line show results for each prize and bars to the right show results collapsed across prizes. Error bars to the right of the dotted line depict ± 1 clustered standard error.




Figure 9. Percentage of decisions for which participants agreed to type “ab” the required amount of times in order to increase the probability of winning a monetary prize, as a function of whether the probability increase was 25% and the prize was small or whether the probability increase was 5% and the prize was large (within-subjects for each potential increase in expected value of the prize draw). Bars to the left of the dotted line show results for each potential increase in the expected value of the prize draw (collapsed across the *high probabilities* and *low probabilities* conditions) and bars to the right show results collapsed across potential increases in the expected value of the prize draw. Error bars to the right of the dotted line depict ± 1 clustered standard error.

General Discussion

How do people decide whether to incur costs to increase their likelihood of success?

In our investigation of this question, we developed a theory called *prospective outcome bias*. According to this theory, people make costly decisions that they expect to feel good about after the associated outcome has been realized. Importantly, people expect costs incurred to increase their likelihood of success to feel less painful when a successful outcome is eventually realized—even when the decision to incur the costs did not cause

the success. Thus, they are most inclined to incur costs to increase their likelihood of success when success is already very likely.

In Study 1, we established that people expect to evaluate their decisions more positively after successes than after failures, even when their decisions had no bearing on the outcome. This finding provides evidence for a core assumption of our *prospective outcome bias* account. In Studies 2–9, we found that people are more motivated to increase high win probabilities than low win probabilities, as predicted by prospective outcome bias. We observed this result in incentive-compatible laboratory tasks (Studies 2, 3, 6, 7, and 9) and in decisions made about real-world scenarios, such as considering the merits of administering a painful treatment to reduce the likelihood of a patient fatality or of hiring an expensive lawyer to increase the likelihood of successfully suing (Study 5 and Study 8). We found that it applies to decisions about whether to incur a cost to increase one's chances of success (Studies 2 and 4–9), as well as decisions about whether to forego a bonus to avoid a decrease in one's chances of success (Study 3). We also found the effect when we randomly sampled the probabilities under consideration (Study 4). Finally, we corroborated an additional prediction of prospective outcome bias: holding the improvement in the expected value of the lottery constant, people are more motivated to slightly improve the chances of obtaining a very good outcome than to substantially improve the chances of a good-but-not-great outcome (Study 9).

In the course of this investigation, we ruled out some other plausible alternative explanations for people's tendency to be more motivated to increase high win probabilities than low win probabilities.

For example, consider that the first prediction of prospective outcome bias—that people will be more motivated to improve their chances of success when those chances are already high—may seem analogous to the goal gradient hypothesis, which asserts that people’s motivation to achieve a goal increases as they get closer to it (Hull, 1932; Kivetz, Urminsky, & Zheng, 2006; Nunes & Drèze, 2006). A leading explanation of goal gradient effects holds that people treat goals as reference points that operate according to the tenets of Prospect Theory (Heath et al., 1999). By this account, people are in the loss domain prior to reaching their goal. Because people are diminishingly sensitive to losses, they are consequently more sensitive to incremental progress toward that goal as they approach it.

Although this account could explain why people are more motivated to increase high chances of success than low chances of success, it cannot explain all of the findings we present here. First, as noted above, the goals-as-reference-point theory predicts that people’s tendency to pay more to *increase* a win probability from 70% to 80% rather than from 20% to 30% would show up just as forcefully when you ask them how much they are willing to pay to enter *separate* lotteries that have win probabilities of 20%, 30%, 70%, and 80%. But the results of Study 7 show that this is not the case. Second, the goals-as-reference-points account does not predict that people’s greater willingness to improve higher chances of success is moderated by variation in outcome bias (Study 8). Finally, it does not predict that, holding constant the improvement in the expected value of their chances of success, people are more motivated to slightly improve the chances of obtaining a very good outcome than to substantially improve the chances of a good-but-not-great outcome (Study 9). To account for this result, a reward that was, say, five times larger

would have to be more than five times as motivating, which would violate the principle of diminishing sensitivity to the magnitude of stimuli (Fechner, 1966). Thus, the goal gradient effect cannot account for the findings in this article.

We also ruled out some other alternative accounts. First, in all of our studies, we ensured that the low and high win probabilities were equally extreme, thus ruling out the possibility that our results could be caused by a symmetric probability weighting function that overweighs small probabilities and (equally) underweighs large probabilities. Furthermore, in Study 7, we showed that any account that invokes greater sensitivity to differences in high versus low probabilities in general (such as an *asymmetric* probability weighting function) cannot explain our findings, either.

Second, in Study 6, we found that the tendency to pay more to increase high versus low win probabilities extends to circumstances in which people know that they will find out whether their decision influenced the outcome. This result rules out an anticipated regret or waste aversion mechanism for our findings, as the likelihood that participants would find out that their decision was regrettable or wasteful was equal across the high versus low win *probability* conditions.

Third, while a version of loss aversion could potentially explain some of the results if specific assumptions are met (i.e., investments of effort and cash are treated as losses unless followed by a positive outcome), loss aversion does not predict the results of Studies 7 and 9 (see Online Supplement 13). Overall, only prospective outcome bias can parsimoniously account for all of the results reported in these nine studies.

Conclusion

We started this project with the belief that people would apply proportional thinking to probabilities, and consequently would be *less* motivated to increase the probability of success when the probability was already high. We were extremely wrong. People are *more* motivated to increase their probability of success when the probability is already very high, a highly robust result that seems to be caused by prospective outcome bias, the tendency for people to make decisions based on how they expect the outcome to make them feel. In this paper, we focused on demonstrating the phenomenon itself, and establishing its cause, but we are excited for future work that explores all of its consequences. We envision many. For example, it is very hard for the Prospect Theory value function to explain the extent of consumers' overspending on insurance (Barseghyan, Molinari, O'Donoghue, & Teitelbaum, 2013; Sydnor, 2010). In contrast, prospective outcome bias predicts overspending most strongly in insurance-like settings. The likely prospect of successfully avoiding a high-stakes, uninsured loss should make overspending on insurance premiums feel justified, even if the probability of using the insurance is very low. We also expect the effects of prospective outcomes bias to extend to many other domains. Medics may be more likely to put patients through (unnecessary) pain to treat conditions that are more versus less likely to improve on their own, and citizens may stand in line to vote for candidates who are already very likely to win, not because they expect their vote to be influential, but because they know that the outcome could make the effort of voting feel worthwhile.

ENDNOTES

1

This equation may seem to make some unlikely predictions, including that people will be willing to invest considerable amounts for tiny improvements in the value of their chances of success. This issue can easily be resolved by defining w as a function of the improvement $\Delta E(S)$ such that, as $\Delta E(S) \rightarrow 0$, $w \rightarrow 0$. This adjustment does not affect any of the predictions we make in this article because all of the predictions apply to cases in which the potential improvement in the expected value of success remains fixed.

2

In Online Supplement 12, we mathematically derive the key predictions from this model.

3

For example, let's assume that the weight placed on the final expected value of the chances of success is 20% (that is, $w = 0.2$). Given the opportunity to improve an already-high 80% chance of winning \$10 to 90%, then the value of the improvement would be $\$1 = (90\% - 80\%) \times \10 , the final expected value of the chances of success would be $\$9 = 90\% \times \10 , and so people would be willing to pay $\$2.60 = (1 - 0.2) \times \$1 + 0.2 \times \$9$. However, given the opportunity to improve a 10% chance of winning \$10 to 20%, the value of the improvement in chances would still be $\$1 = (20\% - 10\%) \times \10 , but the final expected value of the chances of success would now be just $\$2 = 20\% \times \10 , and so people would be willing to pay just $\$1.20 = (1 - 0.2) \times \$1 + 0.2 \times \$2$.

4

For example, given the opportunity to improve their chances of winning \$10 from 80% to 90%, and again assuming $w = 0.2$, people would be willing to pay $\$2.60 = (1 - 0.2) \times \$1 +$

$0.2 \times \$9$. However, given the opportunity to improve their chances of winning a larger reward of \$100 by a smaller margin from 89% to 90%, the value of the improvement in chances would still be $\$1 = (90\% - 89\%) \times \100 , but the final expected value of the chances of success would now be $\$90 = 90\% \times \100 , and so people would be willing to pay much more: $\$18.80 = (1 - .8) \times \$1 + 0.2 \times \$90$.

5

There were a large number of duplicate responses because many participants could not answer an attention check question on their first attempt and so attempted the survey again. As preregistered, we excluded all such participants from our main analysis.

6

In our preregistration, we said that we would collect data from “one Wharton Behavioral Lab session,” which the Wharton Behavioral Lab defines as comprising 12 separate hour-long sessions taking place over the course of a week. In each hour-long session, many participants are run simultaneously.

7

As prospective outcome bias would predict, the effect of greater willingness to improve higher win probabilities increased with participants’ initial motivation to win the prize. In Online Supplement 8, we report this interaction for all of the studies in which we collected an initial measure of motivation.

8

The 20 prizes were: an iPhone, a mountain bike, a trip to Paris, a \$500 Amazon gift card, a mattress, a spa retreat, a MacBook, a 5-year Netflix subscription, an iPad, a widescreen TV, a digital camera, a sofa, a leather office chair, a gold-plated wrist watch, an Amazon

Fire Kindle, a \$500 eBay gift card, a \$600 BestBuy gift card, an \$800 Macy's gift card, a \$500 Walmart gift card, and a \$300 American Express gift card.

9

For this analysis, we dealt with outliers according to our preregistration, by winsorizing participants' willingness-to-pays at the 95th percentile. The complete details, including results with a second preregistered (rank) dependent variable, are in the full write-up of Study 4 in Online Supplement 9.

10

Specifically, the survey software determined winners for each prize draw by randomly drawing an integer from 1 to 100. If this number was less than or equal to the participant's win probability, the participant would win a prize. For example, if the participant increased their win probability from 80% to 90%, then the participant would win a prize if the random integer was less than or equal to 90, but not if the random integer was higher than 90. Thus, the participant would know that their decision to type caused them to win the prize if the random integer was between 81 and 90, and would also know that their decision to type had no impact if the number was less than 81 or greater than 90. As promised in the survey instructions, we informed participants of this random integer when we revealed the result of the prize draw and explicitly told them whether their decision in fact influenced the outcome.

11

The findings presented thus far could also be consistent with loss aversion. In Online Supplement 13, we show (mathematically) that loss aversion makes different predictions

from prospective outcome bias in Study 7 and Study 9, and that it cannot account for what we find in those studies.

12

This example assumes that people do not asymmetrically weigh probabilities. In the presence of both prospective outcome bias and greater sensitivity to differences between high versus low probabilities, we would expect the difference between people's valuations of the 90% and the 80% lotteries to be greater than the difference between their valuations of the 20% and the 10% lotteries. However, we would still expect for that effect to be much smaller than the effect of people's greater willingness to incur costs to improve a win probability from 80% to 90% than from 10% to 20%. This is because, when people are deciding whether to *improve* their chances of success, their greater valuation of the opportunity to improve higher (vs. lower) win probabilities would be driven both by asymmetric probability weighting and by prospective outcome bias. In contrast, when people are deciding whether to *enter* separate lotteries, the greater difference between their valuations of two higher (vs. lower) win probabilities would be driven only by asymmetric probability weighting.

13

We prove that this prediction follows from our model of prospective outcome bias in Online Supplement 12.

14

Specifically, for the selected decision, we randomly generated a bonus payment that the participant would have to spend if they reported being willing to pay at least as much as this bonus payment to either enter the prize draw or increase their win probability. For

example, if the survey software randomly generated a bonus payment of \$0.25, and for the selected prize draw, the participant had reported that they would be willing to pay \$0.40 to increase their win probability from 3% to 9%, then because they were willing to pay more than the bonus payment ($\$0.40 \text{ WTP} > \$0.25 \text{ bonus payment}$), they would not receive the \$0.25, and would instead spend it on entering the prize draw at a 9% win probability. However, if the survey software had instead generated a bonus payment of \$0.50, they would have been willing to pay less than the bonus payment ($\$0.40 \text{ WTP} < \$0.50 \text{ bonus payment}$); thus, they would have received the \$0.50, and entered the prize draw at the baseline 3% win probability. Since participants' reported willingness to pay did not influence the size of the randomly generated potential bonus payment, participants were incentivized to report it honestly.

15

We confirmed the significance of this interaction in an additional preregistered analysis and a further exploratory analysis (see Online Supplement 10).

16

As in Study 6, in four of the eight scenarios, success constitutes attaining a positive outcome (and participants saw the potential probabilities of the positive outcome), and in the other four, success constitutes avoiding a negative outcome (and participants saw the potential probabilities of the negative outcome).

17

After these questions, participants also answered two questions designed to measure their outcome bias with respect to the decision *not* to improve their chances of success in a given scenario. Although we preregistered to incorporate these questions into our overall measure

of outcome bias, we later realized that doing so muddies the interpretation of this measure, making it less a measure of outcome bias and (potentially) more a measure of anticipated regret. When we follow our preregistration and use these items to construct our measure of outcome bias, the results are still very highly significant in our predicted direction. See Online Supplement 11 for details.

18

We then asked participants to imagine that they did *not* add the new service, and asked the same two questions in the same order. To avoid confusing participants, across all scenarios, we always asked the two questions about the decision to act before the questions about the decision *not* to act, and for each possible decision, we asked them the question that assumed the successful outcome before the question that assumed the unsuccessful outcome.

19

In Online Supplement 12, we show (mathematically) that our model of prospective outcome bias makes this prediction, and in Online Supplement 13, we show that loss aversion does *not* make this prediction.

APPENDIX

Links to Preregistration Forms

Study 1: <https://aspredicted.org/47pe5.pdf>

Study 2: <https://aspredicted.org/c4ve6.pdf>

Study 3: <https://aspredicted.org/jw7cs.pdf>

Study 4: <https://aspredicted.org/pu74e.pdf>

Study 5: <https://aspredicted.org/me85x.pdf>

Study 6: <https://aspredicted.org/iz85p.pdf>

Study 7: <https://aspredicted.org/822sn.pdf>

Study 8: <https://aspredicted.org/ek2xh.pdf>

Study 9: <https://aspredicted.org/kb2fw.pdf>

Study S1: <https://aspredicted.org/s7v6k.pdf>

Table of Contents of the Online Supplemental Materials

Table A1 displays the content of the online supplemental materials available at: <https://osf.io/95kem/>.

Table A1
Table of Contents of the Online Supplemental Material

Section	Pages
Online Supplement 1: Exclusions, Attrition, and Reported Sample Sizes	(p. 1)
Online Supplement 2: Study S1	(pp. 2–5)
Online Supplement 3: Study 3 Full Write-up	(pp. 6–10)
Online Supplement 4: Study 4 Full Write-up	(pp. 11–16)
Online Supplement 5: Study 5 Full Write-up	(pp. 17–21)
Online Supplement 6: Comprehension Checks and Attention Checks	(pp. 22–27)
Online Supplement 7: Preregistered OLS Regression Results for Studies 2, 3, 5, 6, and 9	(p. 28)
Online Supplement 8: Interactions With Initial Motivation	(p. 29)
Online Supplement 9: Study 4 Additional Results and Discussion	(pp. 30–31)
Online Supplement 10: Study 7 Analyses With Rank Dependent Variable	(pp. 32–34)
Online Supplement 11: Study 8 Preregistered Analysis	(pp. 35–36)
Online Supplement 12: Mathematical Derivation of Key Predictions	(pp. 37–45)
Online Supplement 13: Mathematical Derivation of the Predictions of Loss Aversion	(pp. 46–68)
References	(p. 69)

BIBLIOGRAPHY

- Arkes, H. R. (1996). The psychology of waste. *Journal of Behavioral Decision Making*, 9, 213–224.
- Baron, J., & Hershey, J. C. (1988). Outcome bias in decision evaluation. *Journal of Personality and Social Psychology*, 54, 569–579. <https://doi.org/10.1037/0022-3514.54.4.569>
- Barseghyan, L., Molinari, F., O'Donoghue, T., & Teitelbaum, J. C. (2013). The nature of risk preferences: Evidence from insurance choices. *American Economic Review*, 103, 2499–2529.
- Becker, G. M., DeGroot, M. H., & Marschak, J. (1964). Measuring utility by a single-response sequential method. *Behavioral Science*, 9, 226–232. <https://doi.org/10.1002/bs.3830090304>
- Caplan, R. A., Posner, K. L., & Cheney, F. W. (1991). Effect of outcome on physician judgments of appropriateness of care. *Journal of the American Medical Association*, 265, 1957–1960. <https://doi.org/10.1001/jama.1991.03460150061024>
- DellaVigna, S., & Pope, D. (2018). What motivates effort? Evidence and expert forecasts. *The Review of Economic Studies*, 85, 1029–1069. <https://doi.org/10.1093/restud/rdx033>
- Delquié, P., & Cillo, A. (2006). Disappointment without prior expectation: A unifying perspective on decision under risk. *Journal of Risk and Uncertainty*, 33, 197–215. <https://doi.org/10.1007/s11166-006-0499-4>
- Fechner, G. (1966). *Elements of psychophysics* (Vol. I). New York, NY: Holt, Rinehart and Winston.

- Fishbach, A., & Finkelstein, S. R. (2012). How feedback influences persistence, disengagement, and change in goal pursuit. In H. Aarts & A. J. Elliot (Eds.), *Goal-directed behavior* (pp. 203–230). New York, NY: Psychology Press.
- Gino, F., Moore, D. A., & Bazerman, M. H. (2009). No harm, no foul: The outcome bias in ethical judgments (SSRN Scholarly Paper No. ID 1099464). Retrieved from Social Science Research Network website: <https://papers.ssrn.com/abstract=1099464>
- Gino, F., Shu, L. L., & Bazerman, M. H. (2010). Nameless + harmless = blameless: When seemingly irrelevant factors influence judgment of (un)ethical behavior. *Organizational Behavior and Human Decision Processes*, 111, 93–101. <https://doi.org/10.1016/j.obhdp.2009.11.001>
- Heath, C., Larrick, R. P., & Wu, G. (1999). Goals as reference points. *Cognitive Psychology*, 38, 79–109. <https://doi.org/10.1006/cogp.1998.0708>
- Hull, C. L. (1932). The goal-gradient hypothesis and maze learning. *Psychological Review*, 39, 25–43. <https://doi.org/10.1037/h0072640>
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47, 263–291. <https://doi.org/10.2307/1914185>
- Kivetz, R., Urminsky, O., & Zheng, Y. (2006). The goal-gradient hypothesis resurrected: Purchase acceleration, illusionary goal progress, and customer retention. *Journal of Marketing Research*, 43, 39–58. <https://doi.org/10.1509/jmkr.43.1.39>
- Lefgren, L., Platt, B., & Price, J. (2015). Sticking with what (barely) worked: A test of outcome bias. *Management Science*, 61, 1121–1136. <https://doi.org/10.1287/mnsc.2014.1966>

- Lipshitz, R. (1989). "Either a medal or a corporal": The effects of success and failure on the evaluation of decision making and decision makers. *Organizational Behavior and Human Decision Processes*, 44, 380–395. [https://doi.org/10.1016/0749-5978\(89\)90015-0](https://doi.org/10.1016/0749-5978(89)90015-0)
- Loomes, G., & Sugden, R. (1982). Regret theory: An alternative theory of rational choice under uncertainty. *Economic Journal*, 92, 805–824. <https://doi.org/10.2307/2232669>
- Marshall, G. W., & Mowen, J. C. (1993). An experimental investigation of the outcome bias in salesperson performance evaluations. *The Journal of Personal Selling & Sales Management*, 13, 31–47.
- Mellers, B. A., & McGraw, A. P. (2001). Anticipated emotions as guides to choice. *Current Directions in Psychological Science*, 10, 210–214. <https://doi.org/10.1111/1467-8721.00151>
- Nunes, J. C., & Drèze, X. (2006). The endowed progress effect: How artificial advancement increases effort. *Journal of Consumer Research*, 32, 504–512. <https://doi.org/10.1086/500480>
- Ratner, R. K., & Herbst, K. C. (2005). When good decisions have bad outcomes: The impact of affect on switching behavior. *Organizational Behavior and Human Decision Processes*, 96, 23–37. <https://doi.org/10.1016/j.obhdp.2004.09.003>
- Sydnor, J. (2010). (Over) insuring modest risks. *American Economic Journal: Applied Economics*, 2, 177–199.
- Tinsley, C. H., Dillon, R. L., & Cronin, M. A. (2012). How near-miss events amplify or attenuate risky decision making. *Management Science*, 58, 1596–1613. <https://doi.org/10.1287/mnsc.1120.1517>

- Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, 211, 453–458. <https://doi.org/10.1126/science.7455683>
- van de Ven, N., & Zeelenberg, M. (2011). Regret aversion and the reluctance to exchange lottery tickets. *Journal of Economic Psychology*, 32, 194–200. <https://doi.org/10.1016/j.joep.2010.11.008>
- Wells, G. L., & Windschitl, P. D. (1999). Stimulus sampling and social psychological experimentation. *Personality and Social Psychology Bulletin*, 25, 1115–1125. <https://doi.org/10.1177/01461672992512005>
- Zeelenberg, M. (1999). Anticipated regret, expected feedback and behavioral decision making. *Journal of Behavioral Decision Making*, 12, 93–106. [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)
- Zeelenberg, M., & Pieters, R. (2004). Consequences of regret aversion in real life: The case of the Dutch postcode lottery. *Organizational Behavior and Human Decision Processes*, 93, 155–168. <https://doi.org/10.1016/j.obhdp.2003.10.001>