PROBABILITY DOMINANCE

Enrico Diecidue, Haim Levy, and Moshe Levy*

Abstract—The most commonly employed paradigms for decision making under risk are expected utility, prospect theory, and regret theory. We examine the simple heuristic of maximizing the probability of being ahead, which in some natural economic situations may be in contradiction to all three of the above fundamental paradigms. We test whether this heuristic, which we call probability dominance (PD), affects decisions under risk. We set up head-to-head situations where all preferences of a given class (expected utility, original or cumulative prospect theory, or regret theory) favor one alternative yet PD favors the other. Our experiments reveal that 49% of subjects' choices are aligned with PD in contradiction to any form of expected utility or prospect theory maximization; 73% are aligned with PD as opposed to preferences under risk aversion and under original and cumulative prospect theory preferences; and 68% to 76% are aligned with PD contradicting preferences under regret theory. We conclude that probability dominance substantially affects choices and should therefore be incorporated into decision-making models. We show that PD has significant economic consequences. The PD heuristic may have evolved through situations of winner-take-all competition.

I. Introduction

EXPECTED utility (EU), prospect theory (PT), and regret theory (RT) are the most popular paradigms for decision making under risk and uncertainty. Together, they explain a large body of empirical evidence in the lab and in the field (Starmer, 2000; Wakker, 2010). Yet in this paper, we describe a simple heuristic with considerable descriptive power: maximizing the probability of being ahead. That is, we introduce the notion of *probability dominance* (PD), under which decision makers choose the prospect with the higher probability of yielding an outcome better than the alternative (in short, coming out ahead or "being ahead") and contrast it with the three main decision-making paradigms mentioned above. We provide strong experimental support for probability dominance, and show that it has substantial economic implications.

Expected utility (von Neumann & Morgenstern, 1947), and prospect theory, which includes both "original" prospect theory (OPT; Kahneman & Tversky, 1979) and "cumulative" prospect theory (CPT; Tversky & Kahneman, 1992), assume that the decision maker evaluates each alternative, or *prospect*, separately and then chooses the prospect with the highest expected utility (or the highest expected PT value). In this standard setting, the order in which the various prospect's outcomes are matched to events is irrelevant. So, for example, if prospect A yields $10 under a given event (state of nature), then it is irrelevant whether, in this particular event, the alternative prospect B yields $5 or $15; only the univariate distributions of each individual prospect matter. Hence these models allow for within-prospect interactions only (Baillon, Bleichrodt, & Cillo, 2015). Regret theory (Bell, 1982; Fishburn, 1982; Loomes & Sugden, 1982) is a bold exception in that the forgone alternative does influence the decision maker; hence, this theory also allows for between-prospects interactions (Baillon et al., 2015). Under regret theory, in fact, the decision maker evaluates a prospect via comparing, state by state, each outcome with the outcome of the forgone prospect. So here the matching of outcomes to events generally matters, even if all events are equally likely.

In this paper, we introduce a simple heuristic that captures the effects of the forgone alternative yet considers between-prospects interactions in a way that differs from regret theory. In particular, we consider the PD heuristic of choosing the alternative that maximizes the probability of being ahead. We evaluate the empirical performance of PD relative to the EU, PT, and regret theory decision models.

We set up situations in which the decision maker faces two prospects, A and B, where A has a higher expected utility or prospect theory value than B (i.e., A is preferred over B by all nondecreasing utility functions under EU or by all PT value functions) but B yields a higher probability of being ahead. Our data show that in this situation, most subjects prefer prospect B over prospect A. This result was obtained in three different experiments and is robust to different experimental setups. It brings to mind Gilboa's (2010) question about the meaning of commonly employed utility functions: "What does the utility function measure, and specifically, does it have anything to do with notions such as well-being or even happiness?" (p. 3). If maximizing the probability of being ahead maximizes one's well-being, then the PD heuristic should be incorporated into economic and financial models despite its contradicting some fundamental axioms of the classical models.

The rest of our paper proceeds as follows. In section II, we provide background and motivation, and section III describes the experiment and our results. Section IV discusses how our findings relate to the literature, and we conclude in section V with a summary and some suggestions for future research.

II. Background and Motivation

Consider monetary prospects A and B with outcomes $x_1^{\mathrm{A}}, \ldots, x_n^{\mathrm{A}}$ and $x_1^{\mathrm{B}}, \ldots, x_n^{\mathrm{B}}$, respectively, under events $s_1, \ldots, s_n$. Regret theory asserts that the preferences are represented as follows:

$$\mathrm{A} \succcurlyeq \mathrm{B} \iff \sum_{i=1}^{n} p_i Q(u(x_i^{\mathrm{A}}) - u(x_i^{\mathrm{B}})) \geq 0. \tag{1}$$

Here the summation is over all possible events; p_i is the probability of event i, u is the real-valued utility function capturing

Received for publication October 31, 2018. Revision accepted for publication July 24, 2019. Editor: Shachar Kariv.

*Diecidue: INSEAD, Decision Sciences; Levy and Levy: Hebrew University.

We are grateful to Shachar Kariv, the editor, and to three anonymous referees for many helpful comments and suggestions. We also thank Jeeva Somasundaram for detailed comments.

A supplemental appendix is available online at http://www.mitpressjournals.org/doi/suppl/10.1162/rest_a_00890.

The Review of Economics and Statistics, December 2020, 102(5): 1006–1020

https://doi.org/10.1162/rest_a_00890

Table 1.—Two Hypothetical Prospects

Ball/Event	Probability	Outcomes Prospect A	Outcomes Prospect B
	1/100	€1	€2
2	1/100	€2	€3
3	1/100	€3	€4
⋮	⋮	⋮	⋮
99	1/100	€99	€100
100	1/100	€100	€1

the utility of outcomes, and Q is a regret function.[1] The function Q captures the decision maker's attitude toward regret by transforming, event by event, the utility differences between the chosen prospect (A in this case) and the forgone prospect (here, B). The regret function is both increasing and antisymmetric; that is, $Q' \geq 0$ and $Q(-z) = -Q(z)$. Note that if Q is linear, then the evaluation in equation (1) is equivalent to expected utility maximization.

Loomes and Sugden (1982) showed that several key experimental findings can be explained by the psychologically appealing idea of regret aversion (RA), which amounts to regret theory preferences with a convex regret function Q. This convexity of Q implies that for any $c > b > a$, we have $Q(c-a) > Q(c-b) + Q(b-a)$. That is, convexity implies that one large, positive, utility difference is preferred over several small differences adding up to the same amount. This implies that the decision maker, when comparing two prospects, is averse to large, negative utility differences.

Many experimental studies support the notion that forgone alternatives do indeed matter (Loomes & Sugden, 1987a, 1987b; Loomes, 1988a, 1988b, 1989; Starmer & Sugden, 1989a, b; Loomes, Starmer, & Sugden, 1991, 1992; Loomes & Taylor, 1992; Starmer, 1992; Bleichrodt, Cillo, & Diecidue, 2010; for a review, see Bleichrodt & Wakke, 2015).[2] Given that most experimental tests of regret theory support regret aversion, we focus on regret aversion.

The simple heuristic we consider, probability dominance, also incorporates the effects of forgone alternatives but models the between-prospect interactions in a different way: choosing the alternative that maximizes the probability of being ahead. To illustrate and motivate this heuristic, consider the two prospects (i.e., event-contingent payoffs) A and B given in table 1.[3] There is an urn with 100 balls numbered from 1 to 100. A single ball is drawn at random, and that ball's number determines the two prospects' outcomes, as shown in the table. Thus, there are 100 possible events, each of which is assumed (for simplicity) to be equally likely. We remark that the univariate distributions of A and B are identical. Thus, EU, PT, and any other univariate theory (i.e., any theory under which prospects are evaluated separately, precluding interactions between prospects) all predict indifference between the two prospects (Savage, 1954). Regret aversion, which does allow for between-prospects interactions, predicts that A is preferred for any increasing utility function and convex regret function.[4] In contrast, prospect B yields a higher probability of being ahead; that is, with a probability of 99%, it yields a better outcome than prospect A. Choosing prospect A yields only a 1% probability of being ahead (its outcome is better only in event 100). Our intuition is that prospect B's higher probability of being ahead will make most individuals prefer it to A. This is the intuition that we set out to explore experimentally.

According to probability dominance, the decision maker (DM) chooses the prospect that maximizes the probability of obtaining a better outcome than the alternative prospect. Formally, B dominates A by PD if and only if

$$\Pr(x_B > x_A) > \Pr(x_A > x_B).^5$$

In table 1's example, prospect B clearly dominates prospect A by PD because

$$\Pr(x_B > x_A) = 0.99 \quad > \quad \Pr(x_A > x_B) = 0.01.$$

Suppose that we now increase the outcome of A in event 100 from €100 to €101. Now A dominates B by first-order stochastic dominance (FSD).[6] Thus, under EU and under PT, all DMs should strictly prefer prospect A—a claim that holds also for all regret-averse decision makers (as they already preferred A when the outcome in event 100 was only €100). The open empirical question is: Will most of these individuals still prefer B?

[1] In the original regret models, the menu of actions is given and must be chosen in accordance with equation (1). Sarver (2008) introduced a broader concept of regret whereby the agent chooses among various menus; he shows that a menu containing fewer options may be preferable to one containing many options because an agent who selects the smaller menu is less likely to make the wrong choice. Thus, Sarver's model differs from both EU and the original RT in that adding more options may be detrimental rather than beneficial.

[2] Levy (2017) derives conditions under which EU, PT, and RT yield the same ranking of prospects.

[3] This example is inspired by Fishburn and LaValle (1988, p. 1224) and Paterson and Diekmann (1988, p. 108).

[4] Formally, for any convex Q we have:

$$Q(u(100)-u(1)) \geq Q(u(100)-u(99)) + Q(u(99)-u(1)) \\ \geq Q(u(100)-u(99)) + Q(u(99)-u(98)) + Q(u(98)-u(1)),$$

which implies

$$Q(u(100)-u(1)) \geq Q(u(100)-u(99)) + Q(u(99)-u(98)) \\ + Q(u(98)-u(97)) + \cdots + Q(u(2)-u(1)),$$

or

$$Q(u(100)-u(1)) + Q(u(99)-u(100)) + Q(u(98)-u(99)) \\ + Q(u(97)-u(98)) + \cdots + Q(u(1)-u(1)) \geq 0.$$

Here the asymmetry of Q is employed in the last step. Thus, prospect A is preferred to B for *any* increasing utility function u and for *any* convex Q.

[5] In the continuous case, this inequality is equivalent to the requirement that $\Pr(x_A > x_B) > 1/2$.

[6] Prospect A dominates prospect B by FSD if and only if $A(x) \leq B(x)$ for all values x and there is at least one strict inequality for some value x_0, where $A(x)$ and $B(x)$ denote the cumulative distributions of the two prospects under consideration. Under EU, this condition implies that all decision makers with nondecreasing utility functions prefer A to B. For a review of FSD, see Levy (2015).

This is the flavor of our experimental setup, although in most cases, we use only six outcomes to simplify the prospects at hand.[7] We set up choices between two prospects where all DMs of a given type (e.g., all EU maximizers, all PT maximizers, all regret averters) should prefer one prospect, while PD implies preference for the other. This head-to-head test allows us to examine whether the PD heuristic affects choices. Indeed, this is exactly what we find. PD fits the subject's choices much better than the three competing models mentioned above. In our main study (see section IIIA), we find that when PD and FSD have opposite predictions, 49% of the choices support PD, in violation of the choice predicted for all individuals with nondecreasing utility functions under EU and for all PT individuals with nondecreasing value utility functions (both OPT and CPT, with and without assuming a probability weighting). When PD and second-order stochastic dominance (SSD)[8] have opposite predictions, 73% of the choices support PD, in violation of the choice predicted for all risk averters under EU (i.e., all EU individuals with concave utility functions), as well as for all PT individuals. When PD and RA have opposite predictions, 76% of the choices support PD, in violation of the choice predicted for all RA individuals. In our robustness studies (studies 2 and 3 described in sections IIIB and IIIC), we obtain similar results.

The head-to-head situations where all DMs of a given class of preferences prefer one prospect while the other prospect dominates by PD, are especially constructed to have this property. This is done in order to enable us to reach unequivocal conclusions about the role of PD. This, of course, does not mean that PD influences choice only in these special situations. Rather, these are the situations that allow us to isolate and gauge the effects of PD. We expect PD to influence choice in general situations, but in these situations, it may be difficult (or even impossible) to differentiate between PD and competing models.

We do not mean to suggest that individuals employ PD as their sole decision-making criterion, and we acknowledge both the prescriptive relevance of EU and the descriptive power of PT and RT. For example, if prospect A described in table 1 were altered so that the outcome in event 100 increases from €100 to €1 million, then we suspect nearly everyone would prefer A to B—even though B still yields a 99% probability of being ahead. We suggest that PD is a natural descriptive factor substantially affecting choice, but it almost certainly operates in concert with other factors.[9] To get an idea about the economic significance of the PD factor, in study 3, we increase the difference in the expected outcomes of the two prospects in favor of the PD inferior prospect. As expected, the bigger this difference is, the smaller the number of choices of the PD dominating prospect. Yet even when the SSD-dominating (and PD inferior) prospect has an expected outcome 20% higher than the PD-dominating prospect, over 30% of the subjects choose according to PD. This indicates that the effect of PD on choice is both strong and economically significant, from which it follows that PD may factor into both asset pricing and the risk premium.

The literature mentions probability dominance in various contexts, where it is known as "the maximum likelihood to be the greatest" (Blyth, 1972), "probability dominance" (Wrather & Yu, 1982; Scarsini, 1985; Sengupta, 1991), maximizing the likelihood of being "better off" (Castagnoli, 1984; Paterson & Diekman, 1988), "the probabilistically prevailing lottery" (Bar-Hillel & Margalit, 1988), and "the most probable winner" (Blavatskyy, 2006). If all outcomes are equally likely, then PD coincides with the "majority rule," which counts the number of events for which $x_B > x_A$ (Birnbaum & Schmidt, 2008; Birnbaum & Diecidue, 2015). In general, however, the majority rule may well be contradicted by PD, which instead focuses on the probability that $x_B > x_A$.[10]

It is somewhat surprising that PD has not received much empirical attention. The reason may be the model's potentially absurd predictions (as when prospect A's outcome in event 100 of table 1 is changed from €100 to €1 million). We aim to rectify that neglect by contributing to the literature as follows:

- Providing substantial experimental evidence supporting PD
- Suggesting that PD operates in concert with standard models, such as EU maximization, thus avoiding absurd predictions
- Quantifying the economic implications of PD

The vast majority of empirical evidence in the lab and in the field is based on binary choices between two-outcome prospects. Such data are appropriately described by the EU,

[7]With six outcomes, we get PD, which is weaker than with 100 outcomes. We chose to have only six outcomes because the number of bounded rationality errors increases with the complexity of the tasks (see Levy, 2008). In study 3, we also have tasks with four and eight outcomes.

[8]A dominates B by SSD if and only if $\int_{-\infty}^{x} [B(t) - A(t)]dt \geq 0$ for all values x and there is a strict inequality for at least one value x_0, where, again, A(x) and B(x) stand for the cumulative distributions of the two prospects under consideration. If this SSD condition is met, it implies that A is preferred to B by all EU maximizers with a concave utility function, that is, for all risk averters under EU. See, for example, Levy (2015).

[9]A simple (but certainly not the only) way to capture this idea is by modeling choice as maximizing the weighted average of two factors: PD and a standard model, such as EU or PT. In online appendix E, we formalize such a model.

[10]If events are not equally likely, it is possible, for example, that prospect A dominates prospect B by the majority rule but prospect B dominates by PD and has a probability of 99% of being ahead. The choices in our experiments were made as simple as possible to avoid errors resulting from misinterpretation of the probabilities presented; that is why we assigned equal probabilities to all events.

It is noteworthy that FSD implies PD when the outcomes of the two prospects are independent. So in that case, there cannot be a situation in which one prospect dominates by FSD and the other by PD. The opposite is not true, however: PD need not imply FSD. Note also that SSD does not imply PD even in the case of independent outcomes (see appendix B for additional details).

TABLE 2.—TASK 1: CONTROL

		Outcomes (€)	
Event	Probability	Prospect A	Prospect B
1	1/6	4.2	3.8
2	1/6	6.3	5.9
3	1/6	12.7	12.6
4	1/6	20.2	18.7
5	1/6	27.6	24.3
6	1/6	37.9	36.5

Prospect A dominates B, in a transparent manner, by first-order stochastic dominance and also by probability dominance. We therefore expect all subjects to choose prospect A.

PT, and RA models. In more complex situations, however, it is likely that other processes and heuristics, which need to be considered and evaluated, will kick in. We introduce PD and evaluate its empirical performance in a moderately complex environment with more than two outcomes.

III. Experiment and Results

Our experiment consists of three studies. Study 1 (186 subjects) is our main study; study 2 (53 subjects) examines the robustness of our results by employing additional tasks with different prospects and different subject populations; and study 3 (45 subjects) examines the boundaries of PD by increasing the magnitude of the outcomes, varying complexity, and increasing the difference between the expected outcomes of the prospects. Overall, six groups of subjects participated in the three studies.

A. *Study 1*

This study comprises five separate choice tasks. In each task, each participant was presented with two prospects and asked to select her preferred prospect. The tasks all involved six equally likely events, and the outcome of each prospect in each event was given in monetary terms. As an example, table 2 summarizes task 1, which served as a control task (see "Tasks and Results" in section IIIA for the results).

As a robustness check, we ran two versions of study 1: a computerized version and a pen-and-paper version; tasks were the same in both versions. The labeling of prospects as A and B was counterbalanced in each version of the experiment. The order of the tasks was randomized in all but two groups (groups 3 and 4). That said, the control task was always presented last. Although the financial incentives are not identical in all versions, the incentive scheme is the same: one of the tasks was chosen at random, and the subject "played" his chosen prospect in this task with real monetary payoffs. Suppose, for example, that prospect A is selected in task 1; then table 2 shows that the highest possible payoff is €37.9 and the lowest possible payoff is €4.2 (see online appendix D for the detailed questionnaire).

Because results were similar across the two versions and across the subject groups, in the main text we report each task's aggregate results. In appendix A, we report the results separately for each version and each subject group.

Subject groups, procedure, and incentives. There were three experimental groups in study 1. We checked for robustness by varying the procedural specifics for each group.

Group 1: Computerized Version, INSEAD-Sorbonne Behavioral Lab. The computerized version took place at the INSEAD-Sorbonne Behavioral Lab (Paris, France) in June 2016 and was entirely computer based (using Qualtrics software). Subjects (students at Sorbonne university) received oral instructions from the experimenter in the lab and went through computer-based written instructions at the start of the experiment. Participants could begin working on their experimental tasks only after answering two practice questions. The order of the experimental tasks was randomized.

The sample size was 54 (average age 23.2, 42% male). Subjects worked on individual computers, and they were randomly allocated into sessions of five subjects who performed the task during the same time period (but at their respective individual paces). Each subject received €4 for participating, and one subject per session was randomly selected to play one of her choices for real (in a randomly chosen task and a random event realization). A participant so selected could earn from €2 to €37.90 extra (the minimum and maximum payoffs across all prospects in all tasks). The average extra payment was €14.60.

Group 2: Pen-and-Paper Version, INSEAD MBAs. The INSEAD experiment took place in the classroom at INSEAD Fontainebleau (France) during September 2016. The 89 subjects (average age 29.5, 69.3% male) were INSEAD MBA students divided into two sections.

Subjects filled out a hard copy questionnaire by pen. The questionnaires were generated by Qualtrics, with randomized question orders as well as counterbalances between prospects A and B, and then printed. The result was ten different versions of the pen-and-paper questionnaire that were randomly allocated to participants. Subjects received oral instructions from the experimenter in the class, and a sample question was explained to them in detail. In addition, written instructions were provided on the questionnaire, and the experimenter was available to answer questions. Two practice questions were also submitted to the participants. One participant made a mistake on the practice questions, and another ten participants did not reply to these questions. These subjects were nonetheless included in the analysis. Excluding these subjects has a negligible effect on the results.

One participant out of ten was selected to play a randomly chosen question for real money. Altogether, nine such subjects were selected; their additional payoff averaged €22.10.

Group 3: Pen-and-Paper Version, Hebrew University Executive MBAs. This pen-and-paper version took place at the Hebrew University of Jerusalem during July 2016. The 45 subjects (average age 34, 68.7% male) were Executive MBA students of two classes.

Subjects filled out a hard copy questionnaire by pen. There were two versions of the questionnaire, with the roles of A and B counterbalanced. Subjects received oral instructions

TABLE 3.—TASK 2: FSD VERSUS PD

		Outcomes (€)	
Event	Probability	Prospect A	Prospect B
1	1/6	20.2	24.3
2	1/6	6.3	12.6
3	1/6	12.7	18.7
4	1/6	37.9	3.8
5	1/6	4.2	5.9
6	1/6	27.6	36.5

Prospect A dominates by FSD, but prospect B dominates by PD because the probability is 83.3% (5/6) that B will yield a higher outcome; 49.4% of the subjects chose prospect B.

from the experimenter in the class, and a sample question was explained to them in detail. In addition, written instructions were provided in the questionnaire, and the experimenter was available to answer any questions.

In this group, all subjects played a randomly selected task with real money. The outcomes were denominated in new Israeli shekels (NIS), for which the exchange rate (at the time of the experiment) was approximately €1 = 5 NIS. Thus, the expected monetary payoff from the tasks was similar to that in the computerized version, where only one subject out of five played. However, the incentives may have been somewhat stronger in this group because all subjects knew that they were playing "for real." The average payment was 12.8 NIS.

Tasks and results. The numbering and order of the tasks described next are for the purpose of a logical presentation. In the experiment itself, the order was randomized except for the control task, which was always the last one.

Task 1: Control Task. The two prospects in this task are those given in table 2, which reveals that prospect A clearly dominates prospect B by FSD. Hence, we expect all subjects who prefer more to less to choose A. Moreover, prospect A yields a higher outcome than B in every event—that is, the probability of being ahead is 100% for prospect A. Because A dominates B by FSD, RA, and PD, we expect all subjects to choose prospect A.

This task was used as a control to gauge the experiment's error level, and it was always the task presented last. This was done to check not only if the subject understood the experimental setup but also that the subjects kept focused until the end of the experiment.

In this control task, the inferior option B was chosen by only 2% of the subjects, a very low error rate. The rest of our analysis excludes the four participants who failed on the control task. Including these subjects yields almost identical results, which is not surprising since they represent but a small fraction of the total sample.

Task 2: PD versus FSD. The two prospects in task 2 are given in table 3. These are exactly the same prospects (in the univariate sense) as those in task 1, but the matching of outcomes to events has changed. Prospect A still dominates by FSD, but now prospect B yields a higher outcome than A in five of the six events; that is, B dominates A by PD with a 83.3% probability of being ahead.

In this task, 49.4% of the subjects chose prospect B, thereby violating first-order stochastic dominance (hence, violating the EU paradigm) while choosing in accordance with probability dominance. Violation of FSD also implies violation of CPT (and of OPT) because all events are equally likely.[11] That is in sharp contrast to the results of task 1, where 98% of the subjects chose prospect A. The difference between the results on these two tasks is highly significant ($p < 0.0001$), even though their univariate distributions are exactly the same. This is a first indication that the matching of outcomes—hence, the outcome of the forgone alternative—strongly affects choice. However, based on this result, we cannot distinguish between probability dominance and regret aversion, which could also affect the results. A direct comparison between PD and RA is conducted in task 5.

One could argue that the difference between the results of tasks 1 and 2 is driven simply by prospect A's dominance being more transparent in task 1.[12] This interpretation is consistent with the notion of probability dominance, which can be viewed as a heuristic for comparing alternatives when neither has an obvious advantage over the other. In tasks 3 and 4, we further analyze the role of PD in such situations.

In appendix B, we prove the following statement: If the outcomes of the two prospects under consideration are statistically independent, then FSD implies PD. In this case, choosing the prospect that is superior by FSD also maximizes PD. It may therefore be tempting to suppose that if the correlation between the outcomes is positive then, a fortiori, FSD $\Rightarrow$ PD. This intuition is misleading, and, in general, FSD $\nRightarrow$ PD even with a positive correlation and despite the result FSD $\Rightarrow$ PD in the case of independent prospects. Yet in the extreme case when the (rank) correlation is +1, we also have the relation FSD $\Rightarrow$ PD. Task 2 illustrates a case where the correlation between prospects is not +1, the prospects are not independent, and the preferences dictated by FSD and PD differ. In this case, choosing by the PD heuristic violates expected utility.

Task 3: PD versus SSD. The two prospects in task 3 are presented in table 4, and each one has the same expected outcome (€16). Prospect A dominates B by second-order stochastic dominance (see panel A of figure 1), which implies that all

[11]In original prospect theory, the DM maximizes $\mathbb{V} = \sum_i \pi(p_i)V(x_i)$; here, the summation is over all possible events, p_i is the probability of event i, $\pi(p_i)$ is the subjective probability weight, x_i denotes the change of wealth in event i, and $V(x)$ is a monotonically increasing value function with $V(0) = 0$, risk aversion for gains ($V''(x) \leq 0$ for $x > 0$), and risk seeking for losses ($V''(x) \geq 0$ for $x < 0$). If all events are equally likely with probability p, then probability weighting simply multiplies the $\mathbb{V}$ of each prospect by the factor $\pi(p)/p$ relative to the $\mathbb{V}$ with the objective probability, p. Thus, for equally likely events, the same preference holds with or without probability weighting. Since $V(x)$ increases monotonically, it follows that if all events are equally likely, then FSD implies dominance for all OPT decision makers (see Levy, 2015). Cumulative prospect theory is specifically designed to avoid violations of FSD even when events are *not* equally likely. Thus, FSD implies dominance for all CPT decision makers in this more general case.

[12]Tversky and Kahneman (1986) and Birnbaum (1997) showed that violations of stochastic dominance can be found in experiments only if the dominance relation is not transparent.

TABLE 4.—TASK 3: SSD VERSUS PD

		Outcomes (€)	
Event	Probability	Prospect A	Prospect B
1	1/6	16	22
2	1/6	4	9
3	1/6	24	34
4	1/6	8	10
5	1/6	32	2
6	1/6	12	19

Prospect A dominates by SSD, so all risk-averse EU maximizers should choose A; this statement also holds for all original PT decision makers and all cumulative PT DMs (see panel B of figure 1). However, prospect B dominates by PD because, with probability 5/6, it yields a higher outcome than A; 72.8% of the subjects chose prospect B.

TABLE 5.—TASK 4: SSD AND PD IN THE SAME DIRECTION

		Outcomes (€)	
Event	Probability	Prospect A	Prospect B
1	1/6	24	22
2	1/6	12	9
3	1/6	8	34
4	1/6	16	10
5	1/6	4	2
6	1/6	32	19

These are the same prospects as in task 3, but the assignment of outcomes to events is different. Prospect A dominates by SSD, OPT, and CPT; it also dominates by PD because, with probability 5/6, it yields a higher outcome than B; 83.1% of the subjects chose prospect A.

FIGURE 1.—CUMULATIVE PROBABILITY DISTRIBUTIONS OF PROSPECTS IN TASK 3

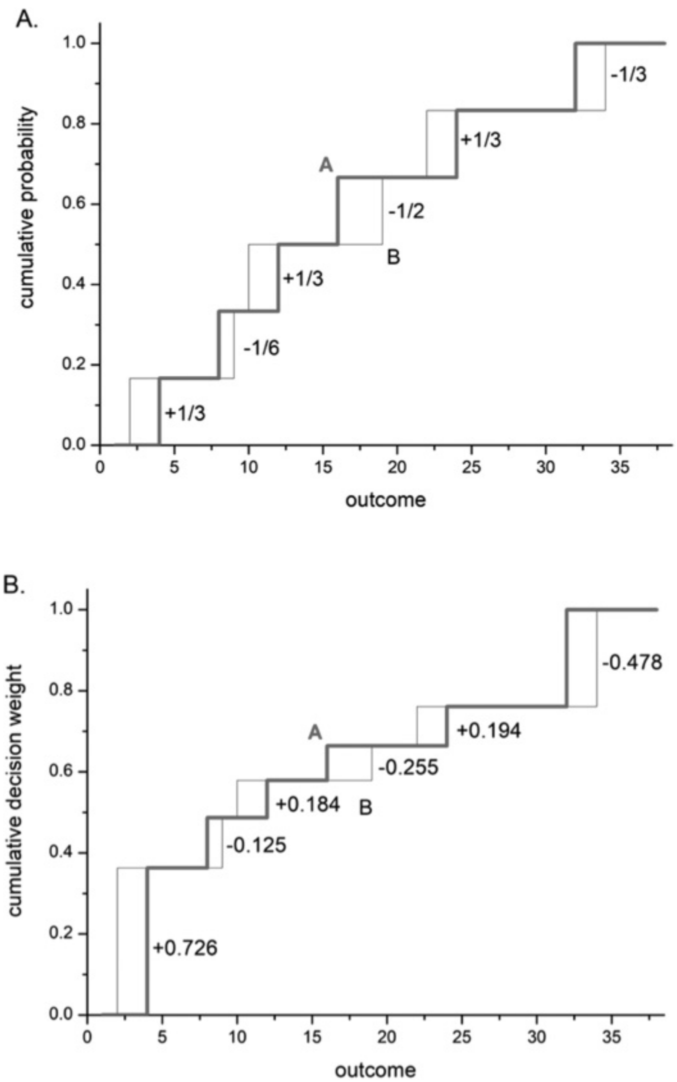

Prospect A (bold line) dominates prospect B (thin line) by SSD because $\int_{-\infty}^{x}(B(t)-A(t))\,dt \geq 0$ for all values of x, where A and B denote the cumulative distributions of prospects A and B, respectively (see Levy, 2015). The numbers in the figure denote the enclosed areas of each rectangle, where a positive number indicates an area in which $B(t) > A(t)$. Panel A is based on objective probabilities of 1/6; panel B uses subjective probability weights with the CPT cumulative probability weighting function and the experimentally estimated value of gamma = 0.61 (Tversky & Kahneman, 1992, equation (6); note that we spell out gamma for the CPT probability weighting function parameter in order to distinguish it from γ, the CRRA parameter employed later). Similarly, prospect A dominates prospect B for any value of the probability weighting function—that is, for any $0 <$ gamma ≤ 1. Thus, all CPT decision makers with any value function and probability weighting function parameters prefer prospect A to prospect B.

risk-averse EU maximizers should prefer prospect A. Since all outcomes are positive and the PT value function is concave in the domain of gains (Kahneman & Tversky, 1979), it follows that the SSD dominance of A also implies that it dominates for any OPT individual, with or without subjective probability weighting of 1/6 (because this is the same probability for all outcomes).[13] Prospect A also dominates prospect B for any DM under CPT for any value function and probability weighting parameters (see panel B of figure 1). By PD, in contrast, prospect B dominates A because the former yields a higher outcome with probability 83.3% (5/6). Note that in this case, the proportion of subjects choosing B is significantly different ($p < 0.01$) between group 1 versus 2 and between group 1 versus 3. Therefore, pooling the three groups is, in this case, not appropriate. We find that 79.5% (55.8%) of the choices favored prospect B in groups 2 and 3 (group 1), in violation of SSD and PT, but in accordance with PD.

Task 4: PD and SSD in the Same Direction. Task 4 pinpoints the importance of PD for choice. The two prospects in this task are exactly identical to the prospects in task 3. The difference here is that the matching of outcomes to events is such that prospect A dominates B not only by SSD, OPT, and CPT but also by PD: A yields a higher outcome than prospect B with probability 5/6 (see table 5). Note that by EU, OPT, CPT, and any other theory based on univariate decision making, the matching of the outcomes should not make any difference, and so the choices made in task 4 should be identical to those in task 3. However, we find that prospect A is chosen by 83.1% of the subjects in task 4 as compared with only 27.2% in task 3. This substantial difference is statistically significant ($p < 0.0001$) and confirms that choices are affected by the outcomes of forgone alternatives.

Task 5: PD versus RA. The evidence from tasks 1 to 4 suggests that univariate models miss an important part of the decision-making process; in particular, many individuals clearly do care about the outcome of the forgone alternative relative to that of the chosen alternative. This property is shared by both regret theory and probability dominance. Regret aversion, which is equivalent to regret theory with a

[13]Because all outcomes are positive, we have $x > 0$ and $V''(x) \leq 0$ in this range. Thus, SSD implies dominance for all OPT decision makers in the case of equally likely events (see note 11 and Levy, 2015).

Table 6.—Task 5: PD and RA Yield Opposite Predictions

Event	Probability	Outcomes (€) Prospect A	Prospect B
1	1/6	5	6
2	1/6	12	14
3	1/6	14	3
4	1/6	9	12
5	1/6	3	5
6	1/6	6	9

Prospect A dominates B for any regret-averse DM, but prospect B dominates prospect A under PD; 76.1% of the subjects chose prospect B.

convex regret function Q, implies that a prospect yielding one large, positive utility difference and several smaller negative differences adding up to the same amount, is preferred over the alternative prospect. This suggests that for a regret-averse DM, prospect A is more attractive in task 3 than it is in task 4: in task 3, A yields one large, positive outcome difference (32 versus 2) and five smaller, negative outcome differences; in task 4, prospect A yields five small, positive differences and one large, negative difference (8 versus 34). This preference is in sharp contrast to subjects' choices in the experiment, where A is chosen by 27.2% of the subjects in task 3 and by 83.1% in task 4. However, this evidence is only suggestive because the regret function Q is defined over differences in utilities and not over differences in outcomes. Task 5 sets up a test of the case where clear-cut opposite predictions are made under probability dominance and regret aversion; the prospects in this task are listed in table 6.

The two prospects in task 5 have the exact same univariate distribution of outcomes. Prospect B clearly dominates A by PD, since it yields a higher outcome than does prospect A in five of the six events. In contrast, B is dominated by A for all regret-averse DMs. To see this, recall that a regret-averse DM prefers A to B if and only if $\sum_{i=1}^{n} p_i Q(u(x_i^{\mathrm{A}}) - u(x_i^{\mathrm{B}})) > 0$ (and prefers B if this value is negative). For the two prospects in task 5, we have

$$\sum_{i=1}^{n} p_i Q(u(x_i^{\mathrm{A}}) - u(x_i^{\mathrm{B}}))$$
$$= \frac{1}{6}[Q(u(5) - u(6)) + Q(u(12) - u(14)) + Q(u(14) - u(3)) + Q(u(9) - u(12)) + Q(u(3) - u(5)) + Q(u(6) - u(9))].$$

Note that the term $Q(u(14) - u(3))$ is positive, while the other five terms are negative. Thus, A is preferred to B if and only if $Q(u(14) - u(3))$ is larger than the absolute value of the five negative terms. By the function's antisymmetric property (i.e., since $Q(-z) = -Q(z)$), we conclude that A is preferred to B if and only if

$$Q(u(14) - u(3)) > Q(u(14) - u(12)) + Q(u(12) - u(9)) + Q(u(9) - u(6)) + Q(u(6) - u(5)) + Q(u(5) - u(3)).$$

This inequality holds for any convex Q because convexity implies that for any $c > b > a$, we have $Q(c - a) > Q(c - b) + Q(b - a)$ (see Loomes & Sugden, 1982; Bleichrodt et al., 2010). Hence for any convex Q, we have

$$\begin{aligned} Q(u(14) - u(3)) &> Q(u(14) - u(12)) + Q(u(12) - u(3)) \\ &> Q(u(14) - u(12)) + Q(u(12) - u(9)) \\ &\quad + Q(u(9) - u(3)) \\ &\vdots \\ &> Q(u(14) - u(12)) + Q(u(12) - u(9)) \\ &\quad + Q(u(9) - u(6)) \\ &\quad + Q(u(6) - u(5)) + Q(u(5) - u(3)). \end{aligned} \tag{2}$$

Thus, A is preferred to B for all regret-averse DMs. In task 5, we find that 76.1% of the subjects chose prospect B—that is, in accordance with PD and in contradiction to RA. This proportion differs significantly from the 50% implied by random choice or indifference ($p < 0.0001$), let alone from the prediction of regret aversion, according to which all decisions should be to choose prospect A.

Note that while the setup designed to contrast PD with RA is very specific, the conclusions from the results in this task are, of course, not limited to this special situation. They indicate that for most subjects, PD provides a better model of choice than RA.

B. *Study 2*

The purpose of study 2 is to examine the robustness of study 1's main findings when different prospects and different subject populations are employed. One of the two subject groups in this study is composed of executives who are, on average, older and more experienced than any of the groups in study 1 (21 subjects); the other subject group consists of students (37 subjects). We find similar results across the two subject groups in this study, so we report only the aggregate results (in appendix A, we report the results for each subject group separately).

Subject groups, procedure, and incentives. There were two experimental groups. As a robustness check we varied the procedural specifics of each group.

Group 4: Hebrew University Executive MBAs. This pen-and-paper version was administered at the Hebrew University of Jerusalem during June 2017. The 21 subjects (average age 38, 78% male) were executive MBA students. Of all our groups, this one contained (on average) the oldest and most experienced subjects.

Participants filled out a hard copy questionnaire by pen. As before, there were two versions of the questionnaire, with the roles of A and B counterbalanced; subjects were given oral instructions by the experimenter in the class, including a sample question explained to them in detail; written instructions were provided in the questionnaire; and the experimenter was available to answer any questions. All subjects played a randomly selected task with real money, with outcomes denominated in NIS. The average payment per subject was 15.6 NIS.

Group 5: Quantitative Students, INSEAD–Sorbonne Behavioral Lab. This experiment took place at the INSEAD–Sorbonne Behavioral Lab (Paris, France) in June 2017 and was also of the pen-and-paper type. The 37 subjects (average age 24.1, 62% male) were economics or science students pursuing bachelor's or master's degrees. Except as explained next, the setup was the same as for group 4. The questionnaires were generated from four different randomizations with a random order of the questions (except for the control question, which was always last). Each subject received €6 for participating, and one subject per session was randomly selected to play one of his or her choices for real. Outcomes were denominated in euros, and the average extra payment was €13.6.

Tasks and results. Study 2 comprises five main tasks,[14] two of which exactly replicate tasks 3 and 5 of study 1 and so are denoted (respectively) tasks 3A and 5A. We employ these tasks to check the robustness of our results across the additional subject populations, including the group of experienced executives. Task 6 is similar to task 3 in this sense: one prospect dominates by SSD and is therefore preferred under EU by any DM with a concave utility function (and is also preferred by all OPT and CPT decision makers); according to PD, though, the other prospect dominates. Yet here the outcomes differ from those employed in task 3. Finally, task 7 is a head-to-head comparison between probability dominance and regret aversion. In this sense it is similar to task 5, but the prospects are different from those in task 5 and have substantially higher outcomes. The control task is once again task 1 of study 1, which is always presented last to check whether subjects remain focused until the end of the experiment. Among the 58 subjects in study 2, the control task was answered correctly by 53 (91%) of them; we report the results for those subjects only (although results are little affected by including all subjects).

Task 3A: PD versus SSD. This is an exact replication of task 3 (as summarized in table 4). Prospect A dominates B by SSD and should be preferred by all DMs with a concave utility under EU, as well as by all OPT and CPT DMs, whereas prospect B dominates by PD. Of the subjects in study 2, 71.7% chose prospect B, thereby supporting the PD model. This percentage is practically the same as in study 1, where 72.8% of the subjects chose prospect B.

Task 5A: PD versus RA. This task replicates task 5 (see table 6); prospect A dominates B for all regret-averse DMs, whereas prospect B dominates A under PD; 69.8% of the subjects chose prospect B, which supports the PD model. The result here is similar to those of task 5 in study 1, where 76.1% of the subjects chose prospect B. Note that the difference between the results of these two studies is not statistically significant ($p = 0.18$). We conclude from the results of tasks 3A and 5A that experienced executives and students with quantitative backgrounds make decisions that are similar to those reported for study 1, and influenced by probability dominance.

Task 6: PD versus SSD. The two prospects in task 6 are given in table 7. As in task 3, prospect A dominates B by SSD (i.e., it is preferred by all concave utility functions under EU). Prospect A also dominates for all OPT and CPT decision makers (see figure 1A in appendix C). In contrast, prospect B dominates A under PD. Observe that the prospects here differ from those employed in task 3, and, in particular, the maximal outcome in task 6 (€42) is larger than its counterpart in task 3 (€34). Prospect B was preferred by 77.4% of the subjects, supporting PD. This percentage is slightly higher than the 72.8% of subjects preferring B in task 3 (although the difference is not statistically significant).

Task 7: PD versus RA. The two prospects in task 7 are given in table 8. As in task 5, prospect A dominates prospect B for all regret-averse DMs whereas A is dominated by B under PD. As compared with the prospects in task 5, however, those here differ and also have larger outcomes: the average outcome in task 7 is €19.7, compared to only €8.2 in task 5. In task 7, 64.2% of the subjects preferred prospect B. This proportion is somewhat lower than the 76.1% of subjects preferring B in task 5, but it is still significantly higher than 50% ($p = 0.019$). The preference for prospect B in this task is also markedly different from the prediction of regret aversion, whereby all subjects should choose prospect A.

[14]The experiment included several additional tasks that were unrelated to the present analysis.

Table 7.—Task 6: SSD versus PD

		Outcomes (€)	
Event	Probability	Prospect A	Prospect B
1	1/6	18	33
2	1/6	35	42
3	1/6	9	15
4	1/6	40	2
5	1/6	28	31
6	1/6	4	11

Prospect A dominates B by SSD, so all risk-averse EU maximizers should choose A, as should all OPT and CPT maximizers (see appendix C). Prospect B dominates prospect A by PD because, with probability 5/6, it yields a higher outcome; 77.4% of the subjects chose prospect B.

Table 8.—Task 7: PD and RA Yield Opposite Predictions

		Outcomes (€)	
Event	Probability	Prospect A	Prospect B
1	1/6	20	25
2	1/6	4	16
3	1/6	16	19
4	1/6	19	20
5	1/6	34	4
6	1/6	25	34

Prospect A dominates prospect B by regret aversion, but B dominates A by probability dominance; 64.2% of the subjects chose prospect B.

It is interesting to note that a large proportion of subjects consistently chose the PD dominating prospect across all the relevant tasks (3A, 5A, 6, and 7): 25 of the 53 subjects in study 2 (47.2%) made these choices. This is significantly higher than the number expected under the assumption of no consistency across tasks, which is only 24.9%.[15] Another 7 subjects (13.2%) chose the PD-dominating prospect in three out of four of these tasks. Only three subjects (5.7%) chose the PD-inferior prospect in all four tasks. Investigating the characteristics of the subjects most influenced by PD is an interesting endeavor left for future work.

C. *Study 3*

The purpose of study 3 is to examine the boundaries of the PD heuristic and its economic significance. As a benchmark, we start with task 3B, which is identical to task 3 in study 1 (and task 3A in study 2), where prospect A dominates by SSD, PT, and CPT, and prospect B dominates by PD. Both prospects have the same expected outcome. In tasks 8, 9, and 10, we gradually increase the difference in expected outcomes in favor of prospect A (as before, the ordering of tasks and labeling of prospects is presented here in this way for convenience; in the experiment they are randomized). Our goal here is to examine the relative roles of the PD heuristic versus the expected outcome in choices. In task 11, we investigate the effect of the magnitude of the outcomes: the prospects are identical to those in task 3 (and 3A and 3B), but with all outcomes multiplied by a factor of 5. Tasks 12 and 13 are designed to examine the role of complexity: in task 12, there are eight equally likely outcomes, and in task 13, there are only four equally likely outcomes.

Subjects, procedure, and incentives. In study 3, there was one group of subjects (labeled group 6). The 45 subjects (average age 27.7, 40% male) were second-year MBA students at Hebrew University. This experiment was conducted as pen-and-paper and was administered at Hebrew University of Jerusalem during March 2019. Subjects were given oral instructions by the experimenter in the class, including a sample question explained to them in detail; written instructions were provided in the questionnaire; and the experimenter was available to answer any questions. All subjects played a randomly selected task with real money, with outcomes denominated in NIS. The average payment per subject was 33.4 NIS.

Tasks and results. Task 3B: PD versus SSD. This is an exact replication of task 3 (as summarized in table 4). Prospect A dominates B by SSD and should be preferred by all DMs with a concave utility under EU, as well as by all OPT and CPT DMs, whereas prospect B dominates by PD. Of the subjects in study 3, 64.4% chose prospect B. Thus, once again we find support for PD. This percentage is somewhat lower than those in studies 1 (72.8%) and 2 (71.7%), perhaps because of some characteristics of the subject population. The main purpose of study 3 is to examine how the support for PD varies in reaction to various factors (difference in expected outcomes, magnitude of the outcomes, complexity), and we use the results of task 3B as a benchmark case. In the following tasks we examine the change in choices of the same subject population in reaction to the above factors.

Tasks 8, 9, and 10: PD versus SSD Increasing Differences in Expected Outcomes. In tasks 8, 9, and 10, prospect A dominates by SSD, while prospect B dominates by PD. As in task 3, prospect B yields a higher outcome with probability 5/6. However, unlike task 3, where the expected outcomes are identical, in tasks 8 to 10, prospect A not only dominates by SSD but also has a higher expected outcome. Table 9 provides these three tasks. In task 8, the expected outcome of prospect A is 6.5% higher than prospect B's expected outcome; in task 9, the difference is 12.9%; and in task 10, it is 19.6%.[16] The percentages of choices of prospect B are 40%,

Table 9.—SSD versus PD: Different Expected Outcomes

		Task 8 Outcomes (NIS)	
Event	Probability	Prospect A	Prospect B
	1/6	16	22
2	1/6	6	8
3	1/6	26	30
4	1/6	8	10
5	1/6	30	4
6	1/6	12	18
		Task 9 Outcomes (NIS)	
Event	Probability	Prospect A	Prospect B
1	1/6	26	34
2	1/6	7	9
3	1/6	18	21
4	1/6	13	18
5	1/6	33	2
6	1/6	8	9
		Task 10 Outcomes (NIS)	
Event	Probability	Prospect A	Prospect B
1	1/6	30	32
2	1/6	6	8
3	1/6	18	20
4	1/6	16	18
5	1/6	34	4
6	1/6	6	10

In all tasks, prospect A dominates prospect B by SSD, but B dominates A by probability dominance (with probability 5/6 of being ahead). In task 8, the expected outcome of prospect A is 6.5% higher than that of B, and 40% of the subjects chose prospect B. In task 9, the expected outcome of prospect A is 12.9% higher than that of B, and 37.8% of the subjects chose prospect B. In task 10, the expected outcome of prospect A is 19.6% higher than that of B, and 31.1% of the subjects chose prospect B.

[15]The percentage of choices of the PD-dominating prospect in these four tasks of study 2 are 71.7%, 69.8%, 77.4%, and 64.2%. Under the assumption of no consistency across tasks, we would expect the percentage of subjects choosing the PD-dominating prospect in all four tasks to be $0.717 \times 0.698 \times 0.774 \times 0.642 = 0.249$, or 24.9%.

[16]In task 8, the expected outcome of prospect B is 15.33, and the expected outcome of prospect A is 16.33: $1/15.33 = 6.5\%$. In task 9, the expected outcome of prospect B is 15.5, and the expected outcome of prospect A is 17.5: $2/15.5 = 12.9\%$. In task 10, the expected outcome of prospect B is 15.33, and the expected outcome of prospect A is 18.33: $3/15.33 = 19.6\%$.

FIGURE 2.—DIFFERENCE IN EXPECTED OUTCOMES AND CHOICES IN TASKS 3B, 8, 9, AND 10

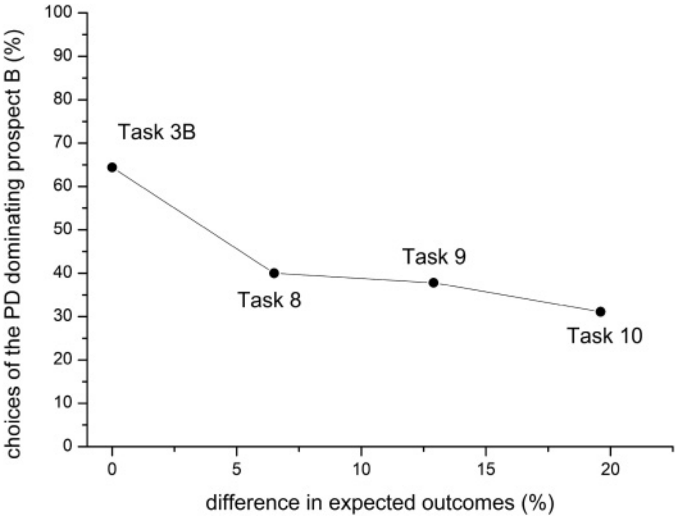

Prospect A dominates B by SSD and has a higher expected outcome than B in tasks 8, 9, and 10. Prospect B dominates A by PD. Even when the expected outcome of A is almost 20% higher than that of B, 31.1% of the subjects preferred the PD-dominating prospect B. The difference in the proportion of subjects choosing prospect B in task 3B and in task 10 is statistically significant at the 1% level.

37.8%, and 31.1%, respectively, shown in figure 2. Thus, as expected, the higher the difference in expected outcomes in favor of prospect A, the higher the proportion of choices in A. The most dramatic decline in the choice of B is from task 3B to task 8. This is reasonable, as some subjects simply choose by the expected outcome. When the expected outcomes are identical, as in task 3B, these subjects resort to other considerations, such as PD. Thus, even a small difference in the expected outcomes makes this subject subpopulation shift to prospect A. Yet even when the expected outcome of prospect A is almost 20% higher than that of B, over 30% of the subjects chose prospect B. These subjects are willing to "pay" 20% of their expected payoff because of the effect of the PD heuristic. This is economically highly significant.[17] In sum, there is no doubt that the expected outcome plays a key role in decision making. Tasks 8 to 10 reveal that PD also plays an important role.

Task 11: Large Outcomes. Task 11 is exactly like the benchmark task 3B, but with all outcomes multiplied by 5. Hence, the expected outcome under both prospects is 80 NIS (compared to 16 NIS in task 3B). In this task, 68.9% of the subjects chose the PD-dominating prospect B (compared with 64.4% in task 3B). The difference is not significant. Yet when the stakes are higher, it seems that the PD factor may become even more important. This suggests that our experimental results may underestimate the importance of PD in real-life situations with high stakes.

[17]This figure is also consistent with an estimate obtained in a different method: in online appendix E, we formalize a model where decisions are made based on a weighted average of standard EU maximization and PD. Based on the results of study 1, we estimate the relative weight attached to the PD factor for various standard utility functions. This relative weight, in turn, implies that subjects are willing to give up as much as 24% in the expected outcome in order to have the PD-dominating prospect.

TABLE 10.—SSD VERSUS PD: VARYING COMPLEXITY

Event	Probability	Task 12: Outcomes (NIS) Prospect A	Prospect B
1	1/8	14	18
2	1/8	28	38
3	1/8	12	10
4	1/8	18	24
5	1/8	6	8
6	1/8	34	4
7	1/8	8	14
8	1/8	16	20
Event	Probability	Task 13: Outcomes (NIS) Prospect A	Prospect B
1	1/4	14	18
2	1/4	22	34
3	1/4	32	12
4	1/4	16	20

In both tasks, prospect A dominates prospect B by SSD, but B dominates A by probability dominance (with probability 3/4 of being ahead). In task 12, 62.2% of the subjects chose prospect B, and in task 13, 66.7% of the subjects chose prospect B. The difference between these proportions is not statistically significant.

Tasks 12 and 13: Varying Complexity. With these tasks, we investigate the role of complexity on the effect of PD. We hypothesize that the PD heuristic becomes more important when the complexity increases. As in the benchmark case, prospect A dominates by SSD, while prospect B dominates by PD (the labels A and B are randomized in the study). The expected outcome is the same in both prospects. In task 12, there are eight equally likely outcomes, and in task 13, there are four equally likely outcomes (compared to the benchmark case with six equally likely outcomes). Tasks 12 and 13 are described in table 10.

We find that the percentage of subjects choosing prospect B is 62.2% in task 12 and 66.7% in task 13 (compared with 64.4% in task 3B). Thus, there is a slight decrease in the choices of the PD-dominating prospect when the complexity increases; however, these differences are not statistically significant. Thus, we do not find a significant effect of complexity, at least when complexity is defined by the number of equally likely outcomes. We did not examine other possible forms of complexity. Other than the complexity of the prospects and the magnitude of their outcomes, the "similarity" between prospects may also play a role determining the importance of PD.[18] We leave the investigation of this issue for future work.

D. *Discussion and Relation to the Previous Literature*

Comparing the results in tasks 1 and 2 and in tasks 3 and 4 clearly reveals that choices are affected by the forgone prospect's outcome. In each pair of tasks, the univariate outcome distributions are the same (i.e., prospect A in task 1 is identical to prospect A in task 2, and so on); the only difference is the assignment of outcomes to events—that is, the "matching" of the outcomes to the prospects. We find that this

[18]We thank one of the anonymous referees for suggesting this point.

Table 11.—PD Consistent with Cycle Observed by Day and Loomes (2010)

		Outcomes		
Event	Probability	Prospect A	Prospect B	Prospect C
1		40	25	15
2	0.05	0	25	15
3	0.05	0	0	15
4	0.80	0	0	0

Day and Loomes (2010) observe the cycle $C \succ A \succ B \succ C$. This intransitivity cycle is consistent with PD because prospect A has a higher probability of being ahead of B (0.10 versus 0.05), prospect B has a higher probability of being ahead of C (0.15 versus 0.05), and prospects A and C each have the same probability of being ahead of the other (0.10).

matching has a significant effect on choice, which contradicts expected utility maximization as well as any other standard theory of univariate decision making (e.g., OPT and CPT). Overall, our evidence confirms the existence of interaction *among* prospects.

The main theory modeling the forgone alternative's influence on choice is regret theory, with the empirically relevant case of regret aversion (RA). When comparing RA with PD (in tasks 5 and 7) in a head-to-head race, we find strong support for probability dominance. How can this finding be reconciled with the large body of empirical research that supports regret aversion? (See Loomes & Sugden, 1987a, 1987b; Loomes, 1988a, 1988b, 1989; Starmer & Sugden, 1989; Loomes, Starmer, & Sugden, 1991, 1992; Loomes & Taylor, 1992; Starmer, 1992; Bleichrodt et al., 2010.)

First, some of the results supporting RA may be driven in large part by event-splitting effects, as Starmer and Sugden (1993) argued; such effects are absent from our experiment. Second, results may be sensitive to the experimental procedures employed (Battalio, Kagel, & Jiranyakul, 1990; Harless, 1992). In this respect, we believe that our choice-based procedure has the advantage of simplicity: there is no eliciting of indifference values, and all events are equally likely. Third, our experiment intentionally focuses on situations with a large difference in the PD between the two prospects, as one prospect yields a higher outcome than the other in five of the six equally likely events. In settings characterized by smaller ΔPD, other effects may prove to be more pronounced. In Bleichrodt et al. (2010), for example, the ΔPD values are typically much smaller (1/5 in two of the cases).

Finally, some of the literature's empirical findings that *contradict* RA are consistent with PD. Consider, for example, the three prospects in table 11 (adapted from Day & Loomes, 2010; Bleichrodt & Wakker, 2015). Day and Loomes empirically observe the intransitive cycle $C \succ A \succ B \succ C$, which is inconsistent with RA (for a discussion of similar cycles, see Tversky 1969).[19] Yet this cycle is consistent with PD because A has a higher probability of being ahead (as compared with B) and B has a higher probability of being ahead (as compared with C). When A and C are compared, each prospect has the same probability of being ahead, and so a preference for C does not violate PD. Some other explanations for intransitivity cycles are based on similarity theory; see, among others, Rubinstein (1988) and Day and Loomes (2010). Under PD, intransitivity cycles are expected to occur even if the prospects are not similar.

[19]In contrast, Birnbaum and Schmidt (2008) and Birnbaum and Diecidue (2015) do not find empirical support for intransitivity cycles.

V. Conclusion

Probability dominance (PD) is a simple heuristic by which a decision maker chooses the prospect with the higher probability of yielding an outcome better than the alternative. This study documents experimentally that PD substantially affects choice. Our experimental design is simple: we set up a choice between two prospects, in which one prospect dominates for *all* preferences of a given class (e.g., under FSD, SSD, OPT, CPT, RA), while the other prospect dominates under PD. In such head-to-head comparisons, we find that (a) 49.4% of the subjects chose according to PD and in contradiction with first-order stochastic dominance, hence, in contradiction to expected utility theory; and (b) 72.8% of the subjects chose according to PD and in contradiction with second-order stochastic dominance (hence, in contradiction to expected utility theory with risk aversion) and in contradiction with prospect theory (both OPT and CPT, and for any parameters). These results indicate that outcomes of forgone alternatives do affect choice, as in regret theory. However, when comparing RA and PD, we find that 68% to 76% of the subjects chose the prospect favored by probability dominance and not the one favored by regret aversion (studies 1 and 2).

Despite the strong experimental support for PD, taken as the sole decision-making criterion, it can lead to absurd choices, as discussed in section II. Thus, from a prescriptive viewpoint, we suggest that PD is one factor influencing choice, together with other factors, such as standard expected utility (or expected value) maximization. We estimate that PD plays an economically significant role in decision making: more than 30% of the subjects were willing to give up a 20% increase in the expected outcome in order to obtain the PD-dominating alternative (study 3).

We expect our PD heuristic to be useful in situations where the alternatives' outcomes are presented event by event rather than as independent univariate distributions. This is the case in many realistic contexts, such as when an investor must choose a mutual fund from a menu of funds and so compares, year by year, the funds' historical returns. Another example is that of a CEO who considers alternative courses of action and conducts a scenario analysis where the outcomes of the alternative actions in each scenario are considered. PD may also explain the phenomenon of investors "picking up pennies in front of a steamroller," behavior that has been suggested as one of the reasons for the 2008 crisis: when comparing a safe asset that yields a certain but low return with an "almost safe" asset that yields higher returns in almost all situations, but also has the potential for a low-probability catastrophic return, in line with PD, many investors may attach more weight to the "almost-always-winning" property of the "almost-safe" asset than to its potential of catastrophe.

What is the origin of the PD heuristic? At this stage, we can only speculate. The PD strategy may be evolutionary advantageous in winner-take-all competitions, where each contestant has to choose one of several possible risky courses of action (or prospects), and the contestant who ends up ahead takes all. In this setting, choosing the PD-dominating prospect maximizes the probability of winning. The winner-take-all setting seems relevant in many evolutionary reproduction competitions. For example, in many animal species, there is a single alpha male who has a large reproductive advantage relative to the other males in the group (see, e.g., Clutton-Brock, 1988).[20]

Expected utility, prospect theory, and regret theory accommodate much of the empirical evidence based on choices between two or three outcome prospects. Yet our research here establishes that PD has a statistically significant and economically substantial effect on choice, especially in situations that are relatively complex. We find consistency in subjects' choices across tasks: subjects who choose according to PD in one task are more likely to choose according to PD in other tasks as well. Future research could investigate the individual characteristics associated with PD and address more precisely the conditions and limitations of the PD heuristic.

Our findings also raise new empirical and experimental questions. Is PD priced in financial markets? Is PD measured with respect to the market index in this context or perhaps relative to other benchmarks? Can PD be simply extended to the choice among more than two alternatives? If so, then does one maximize the probability of obtaining the best outcome among all alternatives or, rather, minimize the probability of obtaining the worst outcome? These are open questions. We believe that the PD heuristic opens new research paths and merits further consideration.

[20]Dekel and Scotchmer (1999) argue that the winner-take-all setting may also induce risk-taking behavior.

REFERENCES

Baillon, A., H. Bleichrodt, and A. Cillo, "A Tailor-Made Test of Intransitive Choice," *Operations Research* 63 (2015), 198–211.

Bar-Hillel, M., and A. Margalit, "How Vicious Are Cycles of Intransitive Choice?" *Theory and Decision* 24 (1988), 119–145.

Battalio, R. C., J. H. Kagel, and K. Jiranyakul, "Testing between Alternative Models of Choice under Uncertainty: Some Initial Results," *Journal of Risk and Uncertainty* 3 (1990), 25–50.

Bell, D. E., "Regret in Decision Making under Uncertainty," *Operations Research* 30 (1982), 961–981.

Birnbaum, Michael H, "Violations of Monotonicity in Judgment and Decision Making" (pp. 73–100), in Anthony A. J. Marley, ed., *Choice, Decision, and Measurement: Essays in Honor of R. Duncan Luce* (Mahwah, NJ: Erlbaum, 1997).

Birnbaum, M. H., and E. Diecidue, "Testing a Class of Models That Includes Majority Rule and Regret Theories: Transitivity, Recycling, and Restricted Branch Independence," *Decision* 2 (2015), 145.

Birnbaum, M. H., and U. Schmidt, "An Experimental Investigation of Violations of Transitivity in Choice under Uncertainty," *Journal of Risk and Uncertainty* 37 (2008), 77.

Blavatskyy, P. R., "Axiomatization of a Preference for Most Probable Winner." *Theory and Decision* 60:1 (2006), 17–33.

Bleichrodt, H., Cillo, A., and E. Diecidue, "A Quantitative Measurement of Regret Theory," *Management Science* 56:1 (2010), 161–175.

Bleichrodt, H., and P. P. Wakker, "Regret Theory: A Bold Alternative to the Alternatives," *Economic Journal* 125 (2015), 493–532.

Blyth, C. R., "Some Probability Paradoxes in Choice from among Random Alternatives," *Journal of the American Statistical Association* 67 (1972), 366–373.

Castagnoli, E., "Some Remarks on Stochastic Dominance," *Rivista di Matematica per le Scienze Economiche e Sociali* 7:1–2 (1984), 15–28.

Clutton-Brock, T. H., ed., *Reproductive Success: Studies of Individual Variation in Contrasting Breeding Systems* (Chicago: University of Chicago Press, 1988).

Day, B., and G. Loomes, "Conflicting Violations of Transitivity and Where They May Lead Us," *Theory and Decision* 68 (2010), 233–242.

Dekel, E., and S. Scotchmer, "On the Evolution of Attitudes towards Risk in Winner-Take-All Games," *Journal of Economic Theory* 87 (1999), 125–143.

Fishburn, P. C., "Nontransitive Measurable Utility," *Journal of Mathematical Psychology* 26:1 (1982), 31–67.

Fishburn, P. C., and I. H. LaValle, "Context-Dependent Choice with Nonlinear and Nontransitive Preferences," *Econometrica* 56 (1988), 1221–1239.

Gilboa, I., "Questions in Decision Theory," *Annual Reviews of Economics* 2 (2010), 1–19.

Harless, D. W., "Actions versus Prospects: The Effect of Problem Representation on Regret," *American Economic Review* 82 (1992), 634–649.

Kahneman, D., and A. Tversky, "Prospect Theory: An Analysis of Decision under Risk," *Econometrica* 47 (1979), 263–291.

Levy, H., "First Degree Stochastic Dominance Violations: Decision Weights and Bounded Rationality," *Economic Journal* 118 (2008), 759–774.

——— *Stochastic Dominance: Investment Decision Making under Uncertainty* (Berlin: Springer, 2015).

——— "Regret Theory: State Dominance and Expected Utility Theory," *Journal of Mathematical Psychology* 79 (2017), 1–12.

Loomes, G., "Further Evidence of the Impact of Regret and Disappointment in Choice under Uncertainty," *Economica* 55 (1988a), 47–62.

Loomes, G., "When Actions Speak Louder Than Prospects," *American Economic Review* 78 (1988b), 463–470.

——— "Predicted Violations of the Invariance Principle in Choice under Uncertainty," *Annals of Operations Research* 19 (1989a).

——— "Experimental Economics" (pp. 152–178), in *Current Issues in Microeconomics* (London: Palgrave Macmillan, 1989B).

Loomes, G., C. Starmer, and R. Sugden, "Observing Violations of Transitivity by Experimental Methods," *Econometrica* 59 (1991), 425–439.

——— "Are Preferences Monotonic? Testing Some Predictions of Regret Theory," *Economica* 59 (1992), 17–33.

Loomes, G., and R. Sugden, "Regret Theory: An Alternative Theory of Rational Choice under Uncertainty," *Economic Journal* 92:368 (1982), 805–824.

——— "Testing for Regret and Disappointment in Choice under Uncertainty," *Economic Journal*, 97 (1987a), 118–129.

——— "Some Implications of a More General Form of Regret Theory," *Journal of Economic Theory* 41 (1987b), 270–287.

Loomes, G., and C. Taylor, "Non-Transitive Preferences over Gains and Losses," *Economic Journal* 102 (1992), 357–365.

Paterson, I., and A. Diekmann, "A Paradox in Decision Theory and Some Experimental Results: The Relative Nature of Decisions," *Theory and Decision* 25 (1988), 107–116.

Rubinstein, A., "Similarity and Decision-Making under Risk (Is There a Utility Theory Resolution to the Allais Paradox?)," *Journal of Economic Theory* 46 (1988), 145–153.

Sarver, T., "Anticipating Regret: Why Fewer Options May Be Better," *Econometrica* 76 (2008), 263–305.

Savage Leonard, J., *The Foundations of Statistics* (New York: Wiley 1954).

Scarsini, M., "A Note on Bernoulli's Principle and Probability Dominance," *Journal of Optimization Theory and Applications* 47:1 (1985), 109–113.

Sengupta, J. K., "Maximum Probability Dominance and Portfolio Theory," *Journal of Optimization Theory and Applications* 71 (1991), 341–357.

Starmer, C., "Testing New Theories of Choice under Uncertainty Using the Common Consequence Effect," *Review of Economic Studies* 59 (1992), 813–830.
——— "Developments in Non-Expected Utility Theory: The Hunt for a Descriptive Theory of Choice under Risk," *Journal of Economic Literature* 38 (2000), 332–382.
Starmer, C., and R. Sugden, "Violations of the Independence Axion in Common Ratio Problems: An Experimental Test of Some Competing Hypotheses," *Annals of Operations Research* 19 (1989a), 79–102.
——— "Probability and Juxtaposition Effects: An Experimental Investigation of the Common Ratio Effect," *Journal of Risk and Uncertainty* 2 (1989a), 159–178.
——— "Testing for Juxtaposition and Event-Splitting Effects," *Journal of Risk and Uncertainty* 6 (1993), 235–254.
Tversky, A., "Intransitivity of Preferences," *Psychological Review* 76:1 (1969), 31.
Tversky, A., and D. Kahneman, "Advances in Prospect Theory: Cumulative Representation of Uncertainty," *Journal of Risk and Uncertainty* 5 (1992), 297–323.
von Neumann, J., and O. Morgenstern, *Theory of Games and Economic Behavior* (Princeton: Princeton University Press, 1947).
Wakker, P. P., *Prospect Theory: For Risk and Ambiguity* (Cambridge: Cambridge University Press, 2010).
Wrather, C., and P. L. Yu, "Probability Dominance in Random Outcomes," *Journal of Optimization Theory and Applications* 36 (1982), 315–334.

Appendix A

Choice by Subject Group

Study 1. Group 1 are students recruited through the INSEAD–Sorbonne Behavioral Lab. Group 2 are INSEAD MBA students. Group 3 are Hebrew University Executive MBA students. In the control task (task 1), 186 of 190 subjects answered correctly, and the results in the table are for these 186 subjects. The asterisks indicate that the differences between group 1 and group 2 (and between group 1 and group 3) are statistically significant at the 1% level (see table A1).

Study 2. Group 4 are Hebrew University Executive MBAs. Group 5 are economics and science students from universities in Paris recruited through the INSEAD–Sorbonne Behavioral Lab. Fifty-three of the 58 subjects answered the control task correctly, and the results in the table are for these 53 subjects (see table A2).

Appendix B

Stochastic Dominance and Probability Dominance: Independent Outcomes

Proposition 1. *If F dominates G by FSD and the outcomes of the two prospects under consideration are independent, then F dominates G also by PD.*

Proof. Let us denote the outcomes of F by x, and the outcomes of G by y, and assume that x and y are continuous and independent. Let $F(x)$ and $G(y)$ denote the respective cumulative distribution functions. The dominance of F over G by FSD implies

$$F(z) \leq G(z) \quad \text{for all } z \text{ (see, for example, Levy 2015).}$$

TABLE A1.—STUDY 1

Task		Group 1: INSEAD-Sorbonne Lab, Computerized (52 subjects) %		Group 2: INSEAD, Pen and Paper (89 subjects) %		Group 3: Hebrew University Pen and Paper (45 subjects) %		Total (186 subjects) %	
		A	B	A	B	A	B	A	B
2	A dominates by FSD, B dominates by PD	51.9	48.1	55.1	44.9	40.0	60.0	50.6	49.4
3	A dominates by SSD, PT, and CPT, B dominates by PD	44.2	55.8*	21.0	79.0*	20.0	80.0*	27.2	72.8
4	A dominates by PD, SSD, and PT (same univariate distributions as in task 3; outcomes are shuffled)	76.9	23.1	86.0	14.0	84.4	15.6	83.1	16.9
5	A dominates by regret aversion, B dominates by PD	23.1	76.9	20.7	79.3	31.1	68.9	23.9	76.1

TABLE A2.—STUDY 2

Task		Group 4: Hebrew University, Pen and Paper (18 subjects) %		Group 5: INSEAD-Sorbonne Lab, Pen and Paper (35 subjects) %		Total (53 subjects) %	
		A	B	A	B	A	B
3A	A dominates by SSD, PT, and CPT; B dominates by PD	11.1	88.9	37.1	62.9	28.3	71.7
5A	A dominates by regret aversion; B dominates by PD	22.2	77.8	34.3	65.7	30.2	69.8
6	A dominates by SSD, PT, and CPT; B dominates by PD;	5.6	94.4	31.4	68.6	22.6	77.4
7	A dominates by regret aversion; B dominates by PD	33.3	66.7	37.1	62.9	35.8	64.2

To prove the PD dominance of F over G, we need to show that $\Pr(x > y) > \Pr(y > x)$. We have:

$$\Pr(x > y) = \int_{-\infty}^{+\infty} dy \int_{x=y}^{+\infty} dx\, h(x, y),$$

where $h(x, y)$ is the joint density function, the first integral with respect to y can take any value, and x is integrated from y up to infinity. Using the independence assumption with $h(x, y) = f(x)g(x)$, we can rewrite:

$$\Pr(x > y) = \int_{-\infty}^{+\infty} \int_{x=y}^{+\infty} f(x)g(y)dxdy$$
$$= \int_{-\infty}^{+\infty} g(y)dy \int_{x=y}^{+\infty} f(x)dx,$$

where $f(x)$ and $g(y)$ are the density functions of x and y, respectively. Thus,

$$A \equiv \Pr(x > y) = \int_{-\infty}^{+\infty} [1 - F(y)]g(y)dy.$$

By symmetry we also have:

$$B \equiv \Pr(y > x) = \int_{-\infty}^{+\infty} [1 - G(x)]f(x)dx.$$

We need to show that $A - B \geq 0$. We have:

$$A - B = \int_{-\infty}^{+\infty} [1 - F(y)]g(y)dy - \int_{-\infty}^{+\infty} [1 - G(x)]f(x)dx.$$

As both x and y are random variables defined over the same range, without loss of generality we can replace both of them with a variable z to obtain,

$$A - B = \int_{-\infty}^{+\infty} [1 - F(z)]g(z)dz - \int_{-\infty}^{+\infty} [1 - G(z)]f(z)dz.$$

Using the assumption that F dominates G by FSD, we have,

$$F(z) \leq G(z) \Rightarrow F(z) + \delta(z) = G(z) \text{ where } \delta(z) \geq 0.$$

Plugging this in the above equation yields,

$$A - B = \int_{-\infty}^{+\infty} [1 - F(z)]g(z)dz$$
$$- \int_{-\infty}^{+\infty} [1 - F(z) - \delta(z)]f(z)dz.$$

Hence,

$$A - B = \int_{-\infty}^{+\infty} F(z)[f(z) - g(z)]dz + \int_{-\infty}^{+\infty} \delta(z)f(z)dz.$$

As $\int_{-\infty}^{+\infty} \delta(z)f(z)dz \geq 0$, we only need to show that

$$\int_{-\infty}^{+\infty} F(z)[f(z) - g(z)]dz \geq 0.$$

Integrating this expression by parts yields

$$\int_{-\infty}^{+\infty} F(z)[f(z) - g(z)]dz = (F(z)[F(z) - G(z)]|_{-\infty}^{+\infty})$$
$$- \int_{-\infty}^{+\infty} f(z)[F(z) - G(z)]dz.$$

And as the first term is equal to 0, we finally obtain

$$\int_{-\infty}^{+\infty} F(z)[f(z) - g(z)]dz = \int_{-\infty}^{+\infty} f(z)[G(z) - F(z)]dz.$$

As F dominates G by FSD we have $F(z) \leq G(z)$ for all z. Therefore, this term is nonnegative. Thus, $\int_{-\infty}^{+\infty} F(z)[f(z) - g(z)]dz \geq 0 \Rightarrow A - B \geq 0$, that is, $\Pr(x > y) > \Pr(y > x)$, which completes the proof. Note that we employ the FSD assumption twice in this proof, and we use the independence of the two distributions.

Proposition 2. *PD does not imply FSD, even when the outcomes of the two prospects under consideration are independent.*

Example: Let F yield an outcome of 1 with probability 1/4 and an outcome of 4 with probability 3/4. Let G yield an outcome of 2 with probability 1/2, and outcome of 3 with probability 1/2. Let the outcomes be independent. Then F dominates G by PD, as $\Pr(x > y) = 3/4$, where x denotes the outcomes of F and y denotes the outcomes of G. Yet it is clear that F does not dominate G by FSD, as the two cumulative distributions cross.

Proposition 3. *When the rank correlation between the two prospects is +1, then FSD ⇒ PD.*

The proof is trivial: if the rank correlation is +1, FSD implies that we have $y_i = x_i + \Delta x_i$ where $\Delta x_i \geq 0$ ($i = 1, 2, \ldots, n$, with at least one strict inequality). This means that $\Pr(y > x) = 1$, and we have PD.

Proposition 4. *SSD does not imply PD, even when the outcomes of the two prospects under consideration are independent.*

Example: Let F yield an outcome of 3.8 with probability 1/2, and outcome of 3.9 with probability 1/2. Let G yield an outcome of 1 with probability 1/4 and an outcome of 4 with probability 3/4. Let the outcomes be independent. Then it is easy to verify that F dominates G by SSD. Yet G dominates F by PD, as $\Pr(y > x) = 3/4$, where x denotes the outcomes of F and y denotes the outcomes of G.

Appendix C

Cumulative Distributions in Task 6

FIGURE 1A.—THE CUMULATIVE PROBABILITY DISTRIBUTIONS OF THE PROSPECTS IN TASK 6

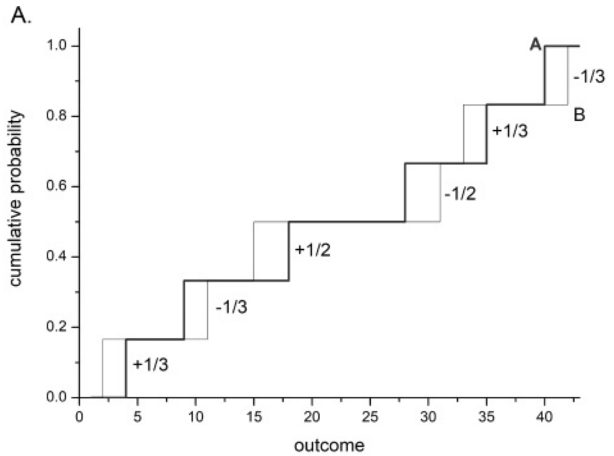

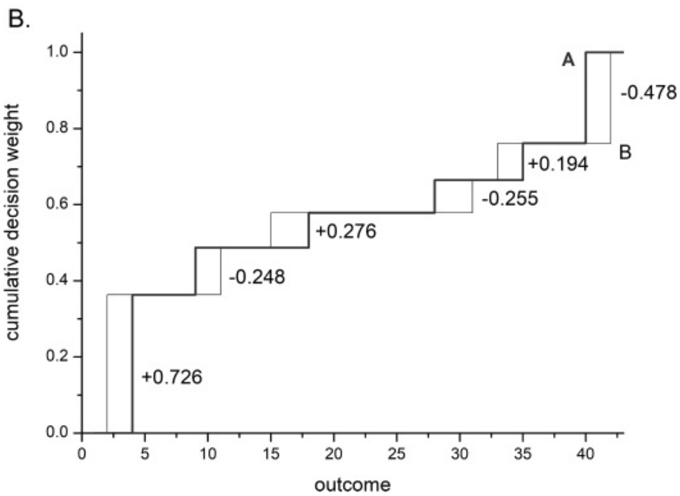

Prospect A (bold line) dominates Prospect B (thin line) by SSD, as $\int_{-\infty}^{x}(B(t) - A(t))dt \geq 0$ for all values of x, where A and B denote the cumulative distributions of A and B, respectively. The numbers in the figure denote the enclosed areas of each rectangle, where a positive number indicates an area where $B > A$. (A) The objective probabilities of 1/6 are employed. (B) The subjective probability weights with the CPT cumulative probability weighting function and the experimentally estimated value of *gamma* = 0.61 (Tversky and Kahneman 1992, eq. 6). Similarly, prospect *A* dominates prospect *B* for any value of the probability weighting function, that is, for any $1 \geq gamma > 0$. Thus, all CPT decision makers with *any* value function and probability weighting function parameters prefer *A* to *B*.