Title
An Investigation of Factors that Influence Resource Allocation Decisions

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 39(0)

Authors
Mistry, Percy K.
Trueblood, Jennifer S.

Publication Date
2017

Peer reviewed

An Investigation of Factors that Influence Resource Allocation Decisions

Percy K. Mistry (pkmistry@uci.edu)
Department of Cognitive Sciences, University of California Irvine, Irvine, CA 92697-5100 USA
Jennifer S. Trueblood (jennifer.s.trueblood@vanderbilt.edu)
Department of Psychology, Vanderbilt University, 111 21st Ave S Nashville, TN 37240 USA

Abstract

We investigate how people allocate a limited set of resources between multiple risky prospects. We found that only a small percentage of decisions followed some form of naive diversification or mean-variance optimization. In general, people were less mean-variance optimal than a naive 1/N heuristic. Aspects of choice sets, such as domain, skew, and second order stochastic dominance, affected resource allocation decisions in a similar manner to their influence on single choice gambles. Individual traits traditionally linked to risk propensity seem to manifest in terms of the degree to which people are inclined to diversify. Lower risk aversion and higher risk seeking traits are linked to increasing diversification. Risk congruency, the degree to which peoples' self-reported and elicited risk aversion matches, moderates how susceptible people are to cost framing nudges. We find evidence for heterogeneous clusters where people either under-weight or over-weight segregated costs, leading to the same nudge producing opposite behavioral results within two risk incongruent groups.

Keywords: resource allocation; risk tolerance; risky choice; individual differences; nudges

Introduction

There are many instances where people have to distribute a limited set of resources between multiple choice options. These can be personal (e.g. investment in a set of retirement funds; constructing a stock portfolio; budgeting household expenses) or institutionalized (e.g. capital allocation, bank lending decisions, government budgets) monetary decisions. These could also be non-monetary decisions such as distribution of labor, time, bandwidth, etc. The choice options often vary in terms of their potential costs and benefits, which are probabilistic in nature. Each choice option may thus be represented as a risky prospect which has some probabilistic distribution of outcomes. There is a large amount of literature that examines the decision making process when people have to select only one out of 2 or more risky prospects, that is, where all resources are invested in a single prospect.

There are limited studies however, that extend this to a resource allocation paradigm, namely, how do people distribute a limited set of resources between 2 or more such risky prospects? Some studies suggest that people follow the 1/N heuristic, which proposes that people tend to naively diversify allocation across the available prospects (Benartzi & Thaler, 2001; Bardolet, Fox, & Lovallo, 2011), although this is often the case only for a subset of the people making these decisions. The normative version of this problem is extensively studied in economics - what optimal strategies *should* people adopt? However there is limited research examining whether people come close to adopting such optimal strategies.

It is important to understand how people deviate from optimality, and to understand what aspects of choice sets influence the resource allocation process. We highlight that measuring optimality and sensitivity to choice sets is more complex in resource allocation tasks compared to simple choice gambles. For simple choice gambles, a cognitive account will typically entail valuation of different prospects, and specification of a deterministic or probabilistic decision rule to compare these valuations. The decision rule for resource allocation needs to be more complex to allow for allocation weights to be placed for each gamble. It needs to take into account aspects such as choice bracketing - whether the valuation of choices is performed at an aggregate portfolio or segregated choice level. These aspects may influence whether drivers of decision making that explain single choice behavior can also explain resource allocation decisions. Further, choice bracketing may also affect how sensitive people are to cost framing nudges, where outcomes are re-framed into a gross higher outcome, set-off by a corresponding cost element. We report an experimental study on resource allocation and show how manipulation of different design factors affects the allocation behavior. The main questions we ask are (1) do people naively diversify?, (2) how sensitive are people to choice set manipulations?, (3) how well do measures of risk traits explain individual differences in resource allocation behavior?, (4) how sensitive are people to cost framing nudges?and (5) how optimal (or sub-optimal) is allocation behavior?

Experiment

In this experiment we test people's preferences for distributing a fixed set of resources between multiple risky prospects. 50 undergraduate students from Vanderbilt University participated in the experiment. The cover story for the task was that participants had to play the role of the head of a company that had the opportunity to invest a fixed amount of money (hypothetical \$100,000) into one or more of 4 possible projects. Participants were advised that all projects had the same expected time to completion and their objective was to maximize the return on the invested amount. They were required to invest all the money, but could distribute this in any proportion between the 4 projects, including allocating no resources to one or more projects. Each project had two possible outcomes - success or failure. They were provided with the probability of success (p_S) and failure (p_F = 1-p_S) for each project, as well as the percentage returns on their investment depending on whether a project succeeded or failed. A successful project always had a positive return, whereas a failed project resulted in either a lower positive or a negative return. The 4 projects always varied in terms of the variability (standard deviation) of return outcomes. Participants were given

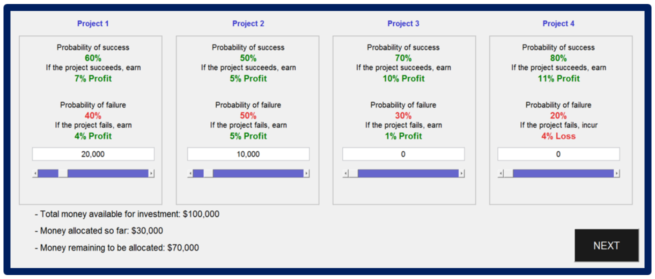

Figure 1: Example interface where participants allocate a fixed set of resources between 4 risky prospects using either text inputs or a moving slider scale.

Figure 2: Example interface of how participants receive feedback for each individual prospect after each trial. The proportion of green to red balls is based on the ratio of probability of success to failure. One of the balls is chosen at random to generate a realized outcome.

an example and a practice trial to familiarize themselves with the interface (see Figure 1). After each trial, participants were provided feedback on the outcome. The outcome was based on the described probability of success and dynamically (randomly) picked by the computer program. The process of realization of the outcome for each project was graphically displayed to the participants. For each project they were shown a box containing 100 X p_S green balls and 100 X p_F red balls. The computer program randomly traversed the box space and eventually picked one of the balls. A green ball implied success and a red ball implied failure (see Figure 2). The returns on the investment for each project were updated based on these outcomes before moving on to the next trial.

(A) Between-subjects conditions

Participants were split into 2 groups of 25 students each. The between subjects design entailed different rewards, with the rest of the design factors being identical between the two groups. Group 1 participated for course credit, and group 2 for financial compensation. This between-subjects condition tests whether financial incentives affect resource allocation behavior. There is mixed evidence for this in tasks involving risky choices (Beattie & Loomes, 1997). Participants in group 2 received a fixed payout of $5 plus an incentive ranging from $0 to $10 that was linked to their performance on the task. For group 2, at the end of the experiment, one of the trials was randomly selected. The incentive component was calculated as $5 plus or minus $0.10 times the %returns achieved on that trial, but limited to the range $0-$10. For example, achieving a loss of 20% resulted in an incentive of $5 - 0.1(20) = $3, and achieving a gain of 20% resulted in an incentive of $5 + 0.1(20) = $7. This allowed for the incentive to be framed as reductions for losses and increments for gains. The total payout including fixed and incentive components ranged between $5-$15.

(B) Within-subjects factorial design

Each participant completed 36 portfolio choice decisions. There were 12 unique decisions based on a 2 (second order stochastic dominance - present vs absent) X 2 (domain - gains vs mixed) X 3 (skew - none, positive, negative) within-subject factorial design. Each of these 12 decisions was repeated in 3 blocks, with the order randomized within blocks. Although the underlying decision remained equivalent across blocks, the three blocks varied in terms of a cost framing effect. The details of the choice set manipulation are given below:

***Second order stochastic dominance (SOSD; 2 levels)*:** In the first level, all prospects in a trial had equal expected value, but the 4 prospects had progressively higher standard deviation. As a result, each prospect had SOSD over the subsequent riskier prospects. In the second level, the prospects were not mean preserving, and riskier prospects (higher standard deviation) also had higher expected values. Thus there was no SOSD. Any behavioral account that is based on a weakly increasing concave utility function, or mean-variance optimization, predicts that prospects with SOSD over other prospects will be a dominated preference. Accordingly, optimal resource allocation under such an assumption would imply allocating 100% of the resources to the safest prospect. SOSD present choice sets allow a parameter free estimation of deviation from mean-variance optimal allocation. On the other hand, choice sets that do not involve SOSD choices allow measuring the level of risk tolerance within a mean-variance optimization framework.

Domain (DM; 2 levels): In the first level, all-gain domain, all outcomes including project failures resulted in positive returns. In the second, mixed domain, the average returns across prospects on *failure* were negative. Domain manipulation allows us to test for the effects of asymmetric gain-loss utilities within the portfolio choice framework.

Skew (SK, 3 levels): In the first level, all prospects had zero skew, that is, success and failure were equally likely. In the second, all except the safest prospect had negative skew, that is, failure outcomes were more likely. In the third, all except the safest prospect had positive skew, that is, success outcomes were more likely. Symmonds, Wright, Bach, and

Dolan (2011) showed that risk and skewness are differently encoded in the brain. People have been shown to be relatively averse to negatively skewed gambles (Deck & Schlesinger, 2010). Manipulation of skew allows us to test whether these effects extend to the portfolio choice paradigm.

Purchase cost framing nudges (PC, 3 levels): In the first level, there are no extraneous purchase costs. In the second and third levels, the outcomes from the first level were translated and re-framed into higher gross outcomes accompanied by an appropriate purchase cost. This re-framing led to prospects that were expected-value-equivalent to the prospects presented in the first level. In the second level, the amount of re-framing was increasingly higher with increasing variability (risk) of prospects outcomes. In the third, the re-framing decreased with increasing variability (risk). The trials were presented in a blocked design with three blocks corresponding to the 3 purchase costs conditions, with the order of the 12 problems in each being randomized. The theory of mental accounting suggests that people may account for re-framed outcomes and corresponding costs in a segregated manner, including under or over weighting the cost component relative to the outcomes. The re-framing thus can act as a nudge, pushing people towards allocation to riskier prospects in the second level and safer in the third level if they underweigh the re-framed costs. The direction of the nudge would be reversed if people over-weigh re-framed costs.

(C) Testing for risk traits

After the allocation task, participants were required to complete one set of paired lottery choices (Holt & Laury, 2002), summarized as HL, with a higher score indicating greater *elicited* risk aversion. They also completed the DOSPERT financial risk-taking (DF) subscale (Blais, 2006), with a higher DF score indicating higher *self-reported* risk seeking behavior, and a locus of control (SL) scale (Rotter, 1966), with higher SL scores indicating a higher *self-reported* external locus of control. All of these influence risky decision making in single choice tasks.

(D) Defining the dependent variables

The simplest way of measuring a resource allocation decision is to look at the *allocation weights* (w_i) for each (i^{th}) prospect, where $\Sigma_{i=1}^{N} w_i = 1$, and N is the total number of choices available. Lopes and Oden (1999) proposed that there are individual differences in whether people approach risky decision making from a perspective of security (protecting low outcomes) or potential (maximizing high outcomes). A simplistic measure of people's security and aspiration levels are measured by the weight allocated to the two extreme prospects - *safest* (S) and *riskiest* (R) respectively. In addition, the *Herfindahl index* $(H) = \Sigma_{i=1}^{N} w_i^2$, where N is the total number of prospects in the choice set, measures the degree of diversification (Rhoades, 1993). When all weights are equal, H takes the minimum value of 1/N (maximum diversification) and when all resources are allocated to a single prospect, H=1. For N=4, values close to 0.25 indicate naive diversification, values close to 0.5 indicate some form of conditional diversification (equal allocation to 2 of 4 prospects), and values close to 1 indicate concentration in a single prospect. These measures S, R, and H reflect segregated measures based on attention paid to individual prospects.

Often, the emergent characteristics of the aggregated portfolio are of greater interest than the individual choices. Most normative theories of portfolio choice are based on optimizing some function of the portfolio characteristics. Since a portfolio can be represented as a probability distribution over outcomes, the most common characteristics are derived from the moments of the resulting portfolio. We calculate the *expected value* (V), and the *standard deviation* (D) of the aggregate portfolio.

Finally, we test for differences between the 2 cost framing conditions. The framing conditions are setup so that correctly accounting for the costs and translation of outcomes should result in no difference between behavior across the three conditions. However, discounting of the costs framed separately would result in a preference for prospects with a higher degree of framing. In one condition, riskier prospects are subject to higher framing (we denote this condition as F_1), and in the other, safer prospects are subject to higher framing, denoted as F_2. Discounting the costs would result in higher selection of riskier prospects in the first and safer in the second framing condition. We calculate *susceptibility to nudges* as, $N = mean[(S_{F_2} - S_{F_1}), (R_{F_1} - R_{F_2})]$. A value of N close to 0 indicates that people are *not* susceptible to cost framing nudges. A high positive value indicates that people underweight separately framed costs, and thus are nudged towards options with higher framing (larger translation of outcomes). A high negative value indicates that people over-weight separately framed costs, and thus are nudged towards options with lower framing (smaller translation of outcomes).

Results

(1) Is there evidence for naive diversification?

Diversification is directly measured using the Herfindahl index (H). The left panel in Figure 3 shows the distribution of H

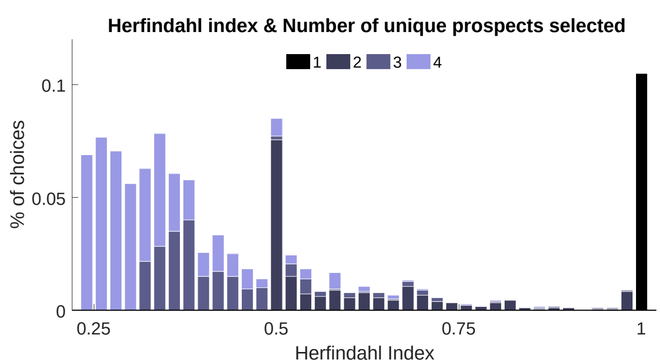

Figure 3: Distribution of Herfindahl index across participants and trials. The color shading shows the number of unique prospects (1 to 4) selected on each trial.

Table 1: Mean values of dependent behavioral measures by design factor. Differences are tested using a Bayesian repeated measures ANOVA for main effects of design factors. Significant differences, measured by log Bayes factors (LBF) ≥ 2.3 are highlighted in bold.

	SOSD		Domain		Skew			Cost framing		
	(Yes)	(No)	(Gain)	(Mixed)	(0)	(Neg)	(Pos)	(None)	(Riskier)	(Safer)
Herfindahl Index (H)	0.48	0.47	0.48	0.47	0.46	0.48	0.49	**0.44**	**0.49**	**0.49**
%Safest (S)	**0.37**	**0.31**	**0.29**	**0.39**	**0.34**	**0.37**	**0.31**	0.32	0.33	0.36
% Riskiest (R)	**0.25**	**0.29**	**0.30**	**0.24**	0.26	0.28	0.27	0.27	0.28	0.26
Expected value (V)	**2.50**	**3.98**	**5.84**	**0.64**	3.23	3.25	3.24	3.27	3.26	3.19
Standard deviation (D)	**3.83**	**4.42**	**2.06**	**6.19**	**4.38**	**3.60**	**4.39**	4.23	4.32	3.83
Susceptibility to nudges (N)	0.01	0.04	0.00	0.05	0.02	0.03	0.02	-	-	-
MPT-error (d_ε)	**0.31**	**0.18**	**0.31**	**0.18**	**0.24**	**0.22**	**0.27**	0.22	0.26	0.25
Risk tolerance (Q_ε)	-	63.7	**19.9**	**107.5**	64.7	52.4	73.9	67.0	76.1	48.0

across participants and trials. The color shading also shows the distribution of the number of unique prospects selected on any trial. A naive diversification strategy would indicate a value of H = 0.25. A large mass of the distribution lies between the range of 0.25 and 0.5, with further peaks at 0.5 and 1.0 indicating choices where people selected 2 of the 4 prospects equally, or invested all their resources in a single prospect, respectively. The 1/N heuristic (naive diversification), proposes that people tend to split allocations evenly between available choices. A variant of this strategy called the conditional 1/N heuristic (Huberman & Jiang, 2006) proposes that people split allocations evenly across a small number of choices rather than the total number of choices available. Using thresholds suggested by Huberman and Jiang (2006), 11% of the choices can be summarized as single prospect concentration, 7% as a conditional diversification into 2 prospects, and 4% as naive diversification into all 4 prospects.

(2) Sensitivity to choice set manipulations:

The mean values of the dependent behavioral variables grouped by experimental factors (which define the type of choice sets) are summarized in Table 1. We conduct a Bayesian repeated measures ANOVA analysis (JASP-Team, 2016) testing for main effects of these design factors. A log Bayes factor, LBF ≥ 2.3 is considered significant, and highlighted in bold in Table 1. There is no evidence that the incentive condition had any effect on S, R, H, V, or D, hence the remainder of the analysis combined data from the course credit and financial incentive conditions. There is no evidence that the level of diversification as measured by H is affected by the domain, skew, or SOSD manipulations. There is evidence for a main effect of domain (LBF 28.1), SOSD (LBF 12.8), and skew (LBF 3.2) on S. Allocation to S is higher in the mixed domain (mean 39%) than in the gains domain (mean 29%), higher in the SOSD (mean 37%) compared to non-SOSD (mean 31%) condition, and higher in negative skew (mean 37%) than positive skew (mean 31%) conditions. There is evidence for a main effect of domain (LBF 9.6) and SOSD (LBF 4.3) on R. Allocation to R is higher in the gains domain (mean 30%) than in the mixed domain (mean 24%), and higher in the non-SOSD (mean 29%) compared to the SOSD (mean 25%) condition. V and D are expected vary with domain and SOSD by design. There is no evidence for a main effect of skew on V, but there is evidence (LBF 6.9) for a main effect of skew on D. Participants exhibit the lowest D (mean 3.6) in the negative skew condition and highest D (mean 4.39) in the positive skew condition, indicating a marked preference for lower variability in the negative skew condition.

(3) Trait-based individual differences

To test if the measured traits influence behavior in the portfolio allocation task, we use a Bayesian ANCOVA analysis treating the between and within subject choice manipulation factors as random effects and testing for the effects of covariates locus of control (SL), risk aversion (HL), and financial risk seeking (DF). We find evidence of an effect of HL (LBF 5.7) and DF (LBF 2.9) on S, and an effect of HL (LBF 13.2) on R. These indicate that higher risk aversion (higher HL and lower DF scores) are linked to higher allocation to the safest prospect and lower allocation to the riskiest prospect, as might be expected. Testing for effects of the locus of control (SL), we find evidence of an effect on S (LBF 3.7), and on R (LBF 6.2). These indicate that increasing external locus of control is also linked to higher allocation to safest and lower to the riskiest prospect. Directionally, this is in contrast to findings based on risky gambles (Rotter, 1966) which showed that increasing external locus of control was associated with waging more money on riskier bets.

Figure 4 shows the joint density of % allocations to R and S. The color coding in the three panels shows the mean level of trait scores. Areas in the centre indicate diversification-like behavior. Interestingly, all 3 mean scores are pretty similar for both extreme decisions (R = 100% and S=100%), but are different for the central areas representing diversification (lower HL scores and higher DF scores). It seems that risk aversion and risk seeking measures are more indicative of how extreme (concentration vs diversification) people are in their allocations, rather than whether they prefer safer or riskier prospects. We find evidence of an effect of HL (LBF 1.4) and DF (LBF 13.4) on H. These indicate that higher risk aversion is linked to lower diversification. This behavior is however contrary to the popular notion that diversification leads to reduced risk.

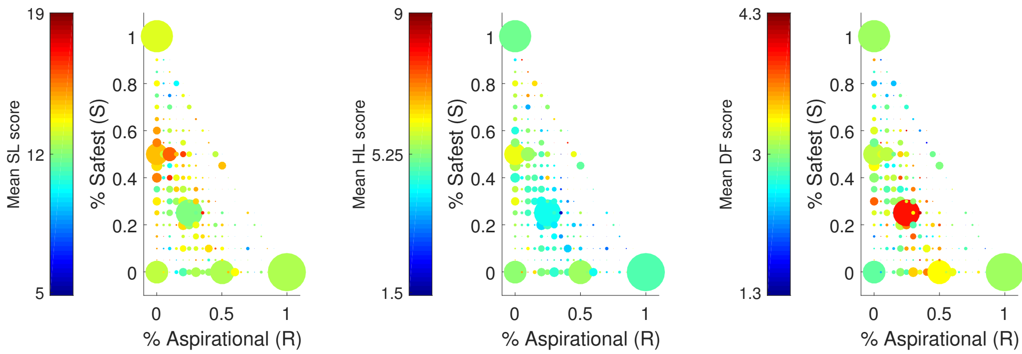

Figure 4: Joint density of % allocation to safest vs aspirational prospects. The size indicates the % of all trials. The 3 panels plot the same density data, with the color coding in the 3 panels showing the mean level fo risk trait scores (*SL*, *HL*, and *DF* respectively).

(4) Sensitivity to cost framing nudges:

Testing for differences between the framing and no framing conditions, we find a significant effect (LBF 10.3) of whether or not there is some cost framing on H, so that framing of either type reduces diversification, with mean H increasing from 0.44 in the no framing condition to 0.49 in both the cost framing conditions. One hypothesis is that the introduction of an additional cognitive element induces people to reduce their diversification. This is supported by the observation that in the no framing condition, people selected all 4 prospects on 55% and either 1 or 2 prospects on 22% of the trials. Compared to this, in the framing conditions (combined), people selected all 4 prospects on 45% and either 1 or 2 prospects on 33% of the trials.

Testing for differences between the two cost framing conditions F_1 (higher risk framing) and F_2 (higher safety framing), a Bayesian repeated measures ANOVA shows evidence that there is no effect on H (LBF -2.7), S (LBF -0.4) or R (LBF -1.9) . Similarly, there is no effect of choice set conditions on N. This seems to indicate that people are not susceptible to differential cost framing effects and adequately account for segregated costs and translation of outcomes. However, we find that risk traits are significant moderators of susceptibility to nudges. We have two measures of risk aversion, a self-reported DF and elicited HL. When these two are congruent, that is, people show both high (low) self-reported and elicited risk aversion, the susceptibility to nudges is lowest, at 0.04. When these are incongruent, and people self-report higher risk aversion but elicited preferences show risk-seeking behavior, the susceptibility to nudges is higher, at 0.06, showing a higher discounting of segregated costs. Most interestingly, when these are incongruent, and people self-report lower risk aversion but elicited preference show higher risk aversion, the susceptibility to nudges is in the reverse direction, at -0.11. This can be interpreted as an over-weighting of segregated costs, leading to a nudge away from choices that had a higher framing effect. The combination of risk congruency and self-reported risk aversion have a significant (LBF 4.7) effect on susceptibility to nudges, and represents a source of significant heterogeneity.

(5) How optimal is allocation behavior?

One of the most popular normative theories of resource allocation is modern portfolio theory (MPT), characterized by mean-variance optimization (Markowitz, 1952). It states that people should select weights that optimize the balance between the expected value and standard deviation of the resulting portfolio. The optimization is a function of a *risk tolerance factor* Q, with lower values of Q indicating preference for safer portfolios and high values of Q indicating preference for riskier portfolios. Given a set of prospects, the theory proposes an *efficient frontier* of possible weight allocations that result in optimization between the desired portfolio mean (V) and variance (standard deviation D). Given an implicit objective to maximize V and minimize D, the frontier represents portfolio choices such that no other combination of weights can result in an increase in V without an increase in D, or a decrease in D without a decrease in V. Note that the efficient frontier does not depend on risk preference Q, but where the selected portfolio lies along the efficient frontier is dependent on the individual preference parameter Q. The set of weights (x) on the efficient frontier for a particular value of Q can be found by minimizing the expression : $x^T \Sigma x - Q E^T x$. Here E is a vector of expected returns on the individual prospects and Σ is the covariance matrix for the returns on the prospects.

Actual portfolios constructed by participants may not lie on the efficient frontier. For any observed portfolio allocation we can calculate the minimum distance of the observed portfolio characteristics from the efficient frontier, which gives the smallest distance to optimality. This distance is dependent on the mean and SD values of individual prospects within a choice set. To enable comparison across choice sets and analyze the impact of factors we calculate the ratio of minimum distance to optimality for the observed portfolio to the largest possible minimum distance to optimality for any combination of weights in the choice set. This is denoted as MPT-error, d_ε. The risk tolerance value corresponding to the closest point on the efficient frontier is denoted Q_ε, and can be inferred to be the risk tolerance level for that choice. Note that in SOSD trials, all prospects have the same EV, and the efficient frontier is a single point that coincides with 100% allocation to the safest prospect. Thus, SOSD trials provide a parameter free estimate of d_ε, but do not allow an estimate for Q_ε.

Table 1 provides the mean MPT-error and risk tolerance levels by design factor levels. Conducting a Bayesian ANOVA analysis and comparing against a null model that included participants as random effects, we find evidence for the effect of SOSD (LBF ∞), domain (LBF ∞), and skew (LBF 2.3) on d_ε. We find evidence for the effect of domain (LBF 36.7) and cost framing (LBF 1.1) on Q_ε. The distance to optimality d_ε measured as a percentage of the largest possible distance to optimality has a mean of (31%; 18%) for SOSD and non-SOSD choice sets; (31%; 18%) for gains and mixed domains; and (24%; 22%; 27%) for no skew, negative skew and positive skew choice sets. As a comparison, d_ε for a 1/N portfolio would be (28%; 13%) for SOSD and non-SOSD sets, (25%; 16%) for gains and mixed domains, and 20% for all skew sets. On an average, the actual allocations that people make are less optimal from a mean-variance optimization standpoint than what a simple 1/N heuristic would result in. The mean d_ε is (22%; 26%; 25%) in the no-framing, and 2 framing conditions. Although the differences are not statistically significant, directionally, framing conditions lead people further away from mean-variance optimality.

The mean inferred risk tolerance Q_ε is 20 in the gains domain and 107 in the mixed domain, indicating that risk tolerance is highly contextual, rather than a stable trait. The mean inferred value of risk tolerance Q_ε is 67 in the no-framing condition, 76 in the higher-riskier-framing condition, and 48 in the higher-safer-framing condition, reflecting sensitivity of risk tolerance to framing effects. Measures of risk traits (SL, HL, DF) do not have any effect on the closest distance to optimality. Evidence is inconclusive (LBF 0.9) for the effect of SL on Q_ε.

Conclusion

The key findings can be summarized as: (1) Only a very small subset of participants follow a naive diversification or 1/N heuristic. (2) Design factors such as domain, skew and SOSD across options influence the allocation that people make in extreme (safest or riskiest) prospects, directionally similar to the effect that these factors have in single choice gambles. (3) We show that individual traits traditionally linked to risk propensity seem to manifest in terms of the degree to which people are inclined to diversify. Lower risk aversion and higher risk seeking traits are linked to increasing diversification. These traits do not seem to be consistently linked to risk tolerance when measured within the MPT framework, and do not seem to influence the relative levels of risk and safety observed in resource allocation behavior. The results are counter-intuitive to the popular notion that diversification is linked to a reduction of risk. (4) We find that cost framing nudges affect the level of diversification. While the effect of nudges seems insignificant at an overall level, a deeper analysis shows trait-based clusters. We find that risk congruency, whether peoples' elicited and self-reported risk aversion are congruent, is a strong moderator for susceptibility to nudges. We find evidence for heterogeneous clusters where people either underweight or over-weight segregated costs leading to the same nudge producing opposite behavioral results in the two risk incongruent groups. (5) We find that people are not optimal under a mean-variance optimization objective, and that on an average, a 1/N heuristic is closer to mean-variance optimization than the actual observed behavior.

References

Bardolet, D., Fox, C. R., & Lovallo, D. (2011). Corporate capital allocation: A behavioral perspective. *Strategic Management Journal*, *32*.

Beattie, J., & Loomes, G. (1997). The impact of incentives upon risky choice experiments. *Journal of Risk and Uncertainty*, *14*.

Benartzi, S., & Thaler, R. H. (2001). Naive diversification strategies in defined contribution saving plans. *American economic review*.

Blais, A.-R. (2006). A domain-specific risk-taking (dospert) scale for adult populations. *Judgment and Decision Making*, *1*(1), 33–47.

Deck, C., & Schlesinger, H. (2010). Exploring higher order risk effects. *The Review of Economic Studies*, *77*.

Holt, C. A., & Laury, S. K. (2002). Risk aversion and incentive effects. *American economic review*, *92*(5).

Huberman, G., & Jiang, W. (2006). Offering versus choice in 401 (k) plans: Equity exposure and number of funds. *The Journal of Finance*, *61*(2), 763–801.

JASP-Team. (2016). Jasp (version 0.8.0.0). *[Computer software]*.

Lopes, L. L., & Oden, G. C. (1999). The role of aspiration level in risky choice: A comparison of cumulative prospect theory and sp/a theory. *Journal of mathematical psychology*, *43*.

Markowitz, H. (1952). Portfolio selection. *The journal of finance*, *7*.

Rhoades, S. A. (1993). The herfindahl-hirschman index. *Fed. Res. Bull.*, *79*, 188.

Rotter, J. B. (1966). Generalized expectancies for internal versus external control of reinforcement. *Psychological monographs: General and applied*, *80*.

Symmonds, M., Wright, N. D., Bach, D. R., & Dolan, R. J. (2011). Deconstructing risk: separable encoding of variance and skewness in the brain. *Neuroimage*, *58*.