

Eyster, Erik; Weizsäcker, Georg

Working Paper

Correlation neglect in financial decision-making

DIW Discussion Papers, No. 1104

Provided in Cooperation with:

German Institute for Economic Research (DIW Berlin)

Suggested Citation: Eyster, Erik; Weizsäcker, Georg (2011) : Correlation neglect in financial decision-making, DIW Discussion Papers, No. 1104, Deutsches Institut für Wirtschaftsforschung (DIW), Berlin

This Version is available at:

<https://hdl.handle.net/10419/52523>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Discussion Papers

1104

Erik Eyster • Georg Weizsäcker

**Correlation Neglect in
Financial Decision-Making**

Berlin, February 2011

Opinions expressed in this paper are those of the author(s) and do not necessarily reflect views of the institute.

IMPRESSUM

© DIW Berlin, 2011

DIW Berlin
German Institute for Economic Research
Mohrenstr. 58
10117 Berlin

Tel. +49 (30) 897 89-0
Fax +49 (30) 897 89-200
<http://www.diw.de>

ISSN print edition 1433-0210
ISSN electronic edition 1619-4535

Papers can be downloaded free of charge from the DIW Berlin website:
<http://www.diw.de/discussionpapers>

Discussion Papers of DIW Berlin are indexed in RePEc and SSRN:
<http://ideas.repec.org/s/diw/diwwpp.html>
<http://www.ssrn.com/link/DIW-Berlin-German-Inst-Econ-Res.html>

Correlation Neglect in Financial Decision-Making

Erik Eyster and Georg Weizsäcker*

December 25, 2010

Abstract

Good decision-making often requires people to perceive and handle a myriad of statistical correlations. Notably, optimal portfolio theory depends upon a sophisticated understanding of the correlation among financial assets. In this paper, we examine people’s understanding of correlation using a sequence of portfolio-allocation problems and find it to be strongly imperfect. Our experiment uses pairs of portfolio-choice problems that have the same asset span—identical sets of attainable returns—and differ only in the assets’ correlation. While any outcome-based theory of choice makes the same prediction across paired problems, subjects behave very differently across pairs. We find evidence for correlation neglect—treating correlated variables as uncorrelated—as well as for a form of “1/N heuristic”—investing half of wealth each of the two available assets. (JEL B49)

Keywords: portfolio choice, correlation neglect, 1/N heuristic, biases in beliefs

*Eyster: Department of Economics, London School of Economics (email: e.eyster@lse.ac.uk); Weizsäcker: DIW Berlin & Department of Economics, University College London & Department of Economics, London School of Economics (email: gweizsaecker@diw.de). We are grateful to Xavier Gabaix, Jacob King, Matthew Rabin, Sven Rady, Tobias Schmidt, Paul Viefers and audiences at Berkeley, DIW Berlin, Bonn, East Anglia, Innsbruck, Princeton, Paris, Toulouse and UCL, as well as the Second Behavioral Economics Annual Meeting (BEAM) at Cornell for their comments and to Jacob King and Wei Min Wang for superb research assistance. Both authors thank the STICERD research centre at LSE for funding the experiments, Technical University Berlin for access to the decision laboratory, and the Schweizerische Nationalbank for funding the research opportunities at the ESSET 2008 conference. Weizsäcker thanks the ELSE Centre at UCL for financial support, and Eyster thanks the NIH for travel support.

1 Introduction

Financial decision-makers face a panoply of correlations across different asset returns. Yet people have limited attention and find it cognitively challenging to work with joint distributions of multiple random variables. Even if the typical investor could in principle adequately analyse financial variables' co-movement, she still might fail to account for correlation properly at the moment of allocating her financial portfolio. As a consequence, investors may hold portfolios that contain undesirable and avoidable risks. Many household investors invest disproportionately in stock of their own employers (Benartzi 2001) or hold only a handful of positively-correlated assets (Polkovnichenko 2005). Discussing the recent financial-markets crisis, numerous commentators have asked whether both households and institutional investors relied on deficient risk modelling.¹

In this paper, we explore people's tendency to ignore correlation. Although such "correlation neglect" may play an important role in numerous economic settings, we focus entirely on its consequences for financial decision-making. The investment behaviours described above are broadly consistent with correlation neglect—but also with a multitude of other forces. To eliminate such confounds, we design and run a series of controlled experiments that test people's attention to correlation. Our experiment studies the standard textbook model of portfolio choice with state-dependent returns and employs a novel framing variation in which each participant faces two versions of the same portfolio-choice problem. Across the two framing variations, we switch asset correlation on and off. Under the null hypothesis that people correctly perceive the covariance structure, the framing variation does not affect their behavior. Nevertheless, we find that behavior changes strongly, and our data analysis supports two alternative hypotheses, both irresponsive to correlation. First, people tend to ignore correlation and treat correlated assets as independent. Second, people tend to follow a simple "1/N heuristic" that prescribes investing equal shares of a financial portfolio into all available assets (as in Benartzi and Thaler (2001)). Our experimental data support these theories despite the fact that, through an understanding test, we ensure that the participants understand the payoff structure, including the co-movement of the asset returns. Consistent with limited attention, participants appear to omit these considerations when choosing

¹Brunnermeier (2009) and Hellwig (2009) discuss erroneous perceptions of systemic risks during the crisis.

their portfolios.

A small set of previous experiments examining people’s responses to correlation uncovers evidence consistent with neglect. Kroll, Levy and Rapoport (1988) and Kallir and Sonsino (2009) find that changing the correlation structure of a portfolio-choice problem leads to little or no change in participants’ decision-making, even when many expected-utility preferences common in the economics and finance literatures predict significant change. Correlation neglect predicts no change in participants’ behavior, consistent with the data.² Our design reverses the previous ones: whereas Kroll et al. (1988) and Kallir and Sonsino (2009) vary the decision problems and find behavior unchanged, we keep the decision problems economically unchanged and find that behavior changes. These two approaches complement each other and together paint a consistent picture of correlation neglect.³

Our experimental design sharpens the findings of this past work because it allows us to test the null hypothesis that people correctly appreciate correlation without making any ancillary assumptions about subjects’ utility functions. Chiefly, we need not assume that subjects are risk averse to test the null hypothesis that they correctly appreciate correlation. In our paired-problems design, any fully rational agent will choose the same distribution of earnings in both investment problems because the set of available portfolios is identical between them. This isomorphy arises because the assets in the correlated frame are linear combinations of those in the uncorrelated frame and therefore span exactly the same set of earnings distributions.⁴ Section 2 presents the main experiment, which involves four such pairs of problems. In two of the problems, the correlated frame uses positive correlation across assets, and in the other two it uses negative correlation. The set of problems with correlated assets thus offers a nontrivial range of hedging opportunities, which may

²Kroll et al. (1988) also report evidence that choices respond remarkably little to feedback on realized returns. Kallir and Sonsino (2009) shed additional light on the cognitive nature of the bias that is consistent with the interpretation of correlation neglect as deriving from limited attention: when asked to predict one asset’s return conditional upon the other’s return, subjects demonstrate that they do perceive correlation correctly despite making investment choices that do not incorporate this understanding.

³Further closely related evidence is provided by Gubaydullina and Spiwoks (2009), whose participants fail to minimize portfolio variance in a problem with correlated assets, but not in a different problem without correlation.

⁴This statement is modulo a qualifier about necessary short-sales constraints, which we clarify in Section 2.

or may not be appreciated by the participants.⁵

Section 3 describes the theoretical predictions for these choice problems, based on a sequence of assumptions. For the standard, “rational” benchmark we first impose the most fundamental assumption that people have rational preferences over portfolios that depend only upon the distribution of monetary payoffs. This consequentialist assumption—encompassing the entire set of expected-utility preferences as well as many generalizations in the literature—implies that if the decision-maker were to make only one choice in our experiment, it would not matter which of the two correlation frames she faced. To invoke this implication in the data analysis, we make a second mild assumption that ensures the decision-maker consider her choices in isolation: we assume that she obeys the Independence Axiom or brackets narrowly between her choices (as defined in Section 3)—making each choice in the experiment without regard to any other. Jointly, these assumptions enable us to reject the null hypothesis that participants correctly perceive correlation, regardless of their risk attitudes. A further (and standard) assumption that we use for some purposes is that the decision-maker is risk averse. Risk aversion suffices to make a unique prediction in the three of our four pairs of problems that have a unique portfolio that second-order stochastically dominates all others. For the remaining pair, we use the stronger assumption that the (risk-averse) decision-maker is close to being risk neutral to make a unique prediction.

The same full set of assumptions suffices to make unique predictions for a decision-maker who fully neglects correlation and treats all random variables as independent. To model correlation neglect, we begin with both assets’ marginal distribution over payoffs and construct their product distribution over payoffs; someone who neglects correlation misperceives payoffs as deriving from this product distribution—where, by construction, the assets’ payoffs are uncorrelated—rather than from the true, correlated joint distribution. Section 3 also introduces the third behavioral model that we consider, a decision-maker who simply invests equal proportions of her wealth in each available asset, the “1/N heuristic”. In the context of our full set of assumptions, we present a model of both covariance neglect and “variance neglect”—ignoring differences in asset variances—

⁵Our design is minimal in these sense that switching correlation between non-degenerate random variables on and off requires at least as many assets (two), states of nature (four), and distinct monetary prizes (three) as we employ.

that nests all three behavioural theories. As described below, the $1/N$ heuristic may be viewed as an extreme version of variance neglect.

The data summary of Section 4 shows that very few participants choose equivalent portfolios in paired choice problems. Only one out of 146 participants in the main experiment behaves fully consistently, choosing four pairs of equivalent portfolios in the four pairs of choices. Of the remaining participants, a majority (60%) does not choose even a single isomorphic pair of portfolios in any of the four pairs of choices. A surprising result appears regarding the relative predictive value of the three basic models (rational, perceived independence, $1/N$). In linear regressions, the rational model adds no explanatory power to the other two: regressing participants' choices on the predictions of the first two or all three models, the rational model garners a point estimate with the wrong sign.

Section 5 further examines the patterns in subjects' deviations from rationality. Because all assets pay the same expected return, our full set of assumptions implies that subjects' preferences over portfolios reduce to preferences over the variance in payoffs.⁶ By allowing people to underestimate both the covariance and variance of asset returns, we build an estimable model that includes all three behavioural models described above in addition to combinations of the three. First, participants' perceptions of correlation may be biased towards zero. Second, they may under-attend to variance; in particular, they may underestimate the magnitude of differences between the entries of the variance-covariance matrix. We model this through a concave transformation of variance. We classify subjects into types according to two parameters, one measuring correlation neglect and the other variance neglect. The single type that best fits the entire subject pool is one that entirely neglects correlation and exponentiates variance to the power 0.43. In a mixture model with type heterogeneity, the two types that best fit the subject pool are one covering ninety-one percent of subjects that essentially coincides with the best-fitting single type just described and a second, rational type best fitting the remaining nine percent of subjects. Adding additional types does not significantly improve the statistical fit. Under the simplifying assumption that all participants either correctly appreciate the covariance structure or show one of the two extreme biases, we

⁶The result in Eyster (2010) clarifies that mild risk aversion leads to approximate mean-variance preferences.

estimate that ten percent of the participants correctly appreciate covariance.

Our study also includes a few more experimental demonstrations of correlation neglect, which we present in Section 6. In one, we offer the participants two portfolios, the first riskier than the second, with the property that if one mistakenly ignores the correlation between the underlying assets the first portfolio appears to first-order stochastically dominate the second portfolio. Indeed, almost all subjects choose the apparently dominant option. But in a control treatment, where the same two portfolios are offered but their true distributions are explicitly shown to the subjects, about half of the subjects choose the second option.

Our final evidence for correlation neglect comes in a demonstration that participants sometimes violate arbitrage freeness. In a separate task we present the participants with three assets, where one asset is state-wise dominated by certain combinations of the other two. Any investment in the dominated asset constitutes an arbitrage loss: spreading that investment appropriately across the other two assets would yield more money in every state of the world. More than three quarters of our subjects fall prey to this arbitrage. Section 7 concludes.

2 Experimental Design and Procedure

The following excerpt from the instructions shows one of the decision problems, labelled Choice 3.

3. Invest each of your 60 points in either Asset E or Asset F, as given below.

	[1]	[2]	[3]	[4]
<i>E</i>	12	24	12	24
<i>F</i>	12	12	24	24

Let $\Omega = \{1, 2, 3, 4\}$ be the four equi-probable possible states of the world, corresponding to columns in the table above. Each row's label X denotes an asset that pays out $X(\omega)$ in state $\omega \in \Omega$. In Choice 3, for instance, Asset E pays out $E(1) = 12$ in state 1, $E(2) = 24$ in state 2, etc. Each entry states how many Euros a participant who invests her entire portfolio (60 "points") in that asset earns. Intermediate investments are rewarded proportionally. The state ω is chosen through a random draw at the end of the experiment. Each participant faces $N = 11$ or $N = 12$ choice problems in this

format, without feedback between choices. Ten problems comprise the first part of the experiment; the remaining one or two choices, described in Section 6, followed after a brief additional instruction. Although the content of this additional instruction varied among participants, its placement after the first ten problems kept it from influencing the initial ten choices. We focus on eight of the initial ten choices in this section, which constitute the main experiment; the remaining two choices are reported in Appendix B.⁷ Only one of the N choices of the experiment is paid out for each participant, following another random draw that each participant makes at the conclusion of the experiment.⁸ Under this random payment procedure, participants can be assumed to make each choice in isolation given preferences that satisfy the Independence Axiom or given that she “brackets narrowly”. We return to this issue in Section 3.

The set of available portfolios in decision problem $n \in \{1, \dots, N\}$ is characterized by the two available assets X_1^n and X_2^n . Any portfolio can be viewed as a lottery over wealth W that assigns wealth $W(\omega)$ to state ω : a decision maker who invests the fraction α_1^n of her wealth ($60 \cdot \alpha_1^n$ “points”) into asset X_1^n and the remaining fraction $1 - \alpha_1^n$ into asset X_2^n ends with wealth $W(\omega) = \alpha_1^n X_1^n(\omega) + (1 - \alpha_1^n) X_2^n(\omega)$ in state ω .⁹ We also require that α_1^n lie in some constraint set, $\mathcal{C}^n \subset \mathbb{R}$, which precludes short sales of either asset, $0 \leq \alpha_1^n \leq 1$; in some cases it also includes more stringent constraints. The following table shows our main eight choice problems by reproducing the eight specifications of $(X_1^n, X_2^n; \mathcal{C}^n)$ for $n = 1, \dots, 8$. The experiment presents the problems in two distinct sequences, attaches abstract labels to states, and varies the order of presentation of states.

⁷The two remaining choices exactly resemble those discussed in the main text and were designed to demonstrate that correlation neglect can lead to violations of blatantly obvious state-wise dominance between assets. Appendix B reports a statistically significant (but economically small) effect of this nature.

⁸In addition, participants receive a show-up fee of 5 Euros.

⁹Although the participants had to choose integer point allocations, placing α_1^n onto a grid, for simplicity we ignore this complication.

n	$\{X_1^n(1), X_1^n(2), X_1^n(3), X_1^n(4)\}\{X_2^n(1), X_2^n(2), X_2^n(3), X_2^n(4)\}$	\mathcal{C}^n
1	$\{15, 21, 15, 21\}\{12, 12, 24, 24\}$	$0 \leq \alpha_1^1 \leq 1$
2	$\{18, 30, 6, 18\}\{12, 12, 24, 24\}$	$0 \leq \alpha_1^2 \leq \frac{1}{2}$
3	$\{12, 24, 12, 24\}\{12, 12, 24, 24\}$	$0 \leq \alpha_1^3 \leq 1$
4	$\{12, 24, 12, 24\}\{12, 18, 18, 24\}$	$0 \leq \alpha_1^4 \leq 1$
5	$\{14, 21, 14, 21\}\{14, 14, 21, 21\}$	$0 \leq \alpha_1^5 \leq 1$
6	$\{14, 21, 14, 21\}\{14, 0, 35, 21\}$	$\frac{2}{3} \leq \alpha_1^6 \leq 1$
7	$\{12, 30, 12, 30\}\{12, 12, 30, 30\}$	$0 \leq \alpha_1^7 \leq 1$
8	$\{12, 30, 12, 30\}\{12, 18, 24, 30\}$	$0 \leq \alpha_1^8 \leq 1$

Table 1: 8 portfolio-choice problems.

An important feature of the eight problems is that half—those with odd-numbered n —involve only pairs of uncorrelated assets, whereas the other half involve non-zero correlations. In Section 3, we explain how each even-numbered decision problem is isomorphic to its immediate predecessor in the table.

The participants for the eight decision problems described above were 148 students of Technical University Berlin, mostly undergraduate. Of these, two participants violated one of the constraints \mathcal{C}^n and their data were excluded from the analysis, leaving 146 complete observations. The eight experimental sessions were conducted in a paper-and-pencil format (with a German translation of the instructions in Appendix A) following a fixed protocol and with the same experimenters present. Each session lasted about 90 minutes, including all payments. Before the main part of the experiment, the participants underwent an understanding test asking three questions about the payment rule, all of which the participants needed to answer correctly before proceeding. About ten percent of participants needed help from the experimenters to pass the understanding test.¹⁰ Controlled variations of the eight decision problems were also used in four further sessions with 96 additional participants (see Section 6).

¹⁰When aiding participants, the experimenters did not produce the correct answer to ensure that each participant find it independently.

3 Predictions

In this section, we develop predictions in our experiments that derive from various different assumptions about the decision maker’s choice behaviour, moving from weakest to strongest. Let

$$Z^n = \{z \in \mathbb{R} : z = \alpha_1^n X_1^n(\omega) + (1 - \alpha_1^n) X_2^n(\omega'), \omega, \omega' \in \Omega, \alpha_1^n \in \mathcal{C}^n\},$$

which if it included the restriction that $\omega = \omega'$ would be the set of feasible money payoffs in choice n of the experiment.¹¹

Without that restriction, Z includes money payoffs only possible if the state governing the payoff of the first asset differed from that governing the payoff of the second, i.e., if the two assets’ correlation was different from what it is. We shall take advantage of this flexibility below in our discussion of agents who ignore correlation. For a rational agent, however, we do impose the restriction that $\omega = \omega'$ and let $\mathcal{W}^n = \mathcal{W}(X_1^n, X_2^n; \mathcal{C}^n)$ be the set of payoff lotteries feasible in decision problem n in isolation—a subset of lotteries over Z^n . Further, let $Z = \cup_n Z^n$, and $\mathcal{W} = \cup_n \mathcal{W}^n$. Because participants’ payments depend upon a single choice chosen at random from across the entire experiment, we work with their preferences over $\Delta\mathcal{W}$, the set of probability distributions over \mathcal{W} . For any $S \subset \Delta Z$, define $m_{\succeq}(S) = \{s \in S : \forall s' \in S, s \succeq s'\}$, the investor’s set of preferred lotteries from S .

Assumption A. *The investor makes choices that maximise a preference relation \succeq over $\Delta\mathcal{W}$ that is complete and transitive. Moreover, for each n , $m_{\succeq}(\mathcal{W}^n)$ is a singleton set.*

Assumption A implies that investors’ preferences over asset portfolios are rational; they depend only on the span of assets and not the correlation structure underlying it. It also assumes a unique optimum. When different assets pay different expected returns, the risk-return tradeoff normally produces non-singleton indifference curves. In our experiment, however, each asset pays the same expected return, which eliminates any risk-return tradeoff. It therefore makes sense to assume that indifference curves are degenerate in each problem. Importantly, our main statistical test requires

¹¹Because subjects invested sixty indivisible tokens, formally the α^n ’s lie on a grid, rendering each Z^n a finite set. We ignore the grid in formal definitions and results to come yet define probability distributions over the Z^n without measurability restrictions because of the finiteness of these sets.

merely that the decision maker have a unique preference-maximising portfolio in each choice set included in the experiment. This implies that whenever two portfolio-allocation problems have the same asset span, the decision maker must make the same choice in each. In particular, a participant confronted with only a single choice problem in the experiment would make a choice that depended only upon the span of assets in that problem. This observation holds not only for expected-utility preferences but for all rational preference relations, regardless of whether they are continuous or satisfy the Independence Axiom.

Observation 1. *Under Assumption A, if $\mathcal{W}(X'_1, X'_2; \mathcal{C}') = \mathcal{W}(X_1, X_2; \mathcal{C})$, then $m_{\succeq}(\mathcal{W}(X'_1, X'_2; \mathcal{C}')) = m_{\succeq}(\mathcal{W}(X_1, X_2; \mathcal{C}))$.*

In our experiment, twinned problems are constructed to have identical asset spans. For example, in Choices 3 (described in the last section) and Choice 4 (below), Asset G is identical to Asset E, and Asset H equals $\frac{1}{2}E + \frac{1}{2}F$.

4. Invest each of your 60 points in either Asset G or Asset H, as given below.

	[1]	[2]	[3]	[4]
<i>G</i>	12	24	12	24
<i>H</i>	12	18	18	24

Someone who invests $\alpha \geq \frac{1}{2}$ in Asset E of Choice 3 can achieve the same portfolio with $\hat{\alpha} = 2\alpha - 1$ in Asset G of Choice 4. This holds because the remaining $1 - \hat{\alpha} = 1 - (2\alpha - 1) = 2(1 - \alpha)$ invested in Asset H, itself comprised of $\frac{1}{2}E + \frac{1}{2}F$, gives $\frac{1}{2}2(1 - \alpha) = (1 - \alpha)$ invested in Asset F, just like in Choice 3. Hence, any portfolio produced using $\alpha \geq \frac{1}{2}$ in Choice 3 can be reproduced through a suitably constructed portfolio in Choice 4, and *vice versa*. Although this argument breaks down for $\alpha < \frac{1}{2}$, the symmetry of Choice 3 across assets and states suggests that anyone who wishes to invest less than one-half of her portfolio in Asset E, essentially betting on State 3 over the symmetric State 2, should be equally willing to invest more than one-half of her portfolio in Asset F. The other three pairs of twinned problems have been constructed similarly so that every feasible portfolio in the uncorrelated problem can be replicated in the correlated problem, and *vice versa*.

Observation 1 implies that if each subject in our experiment were to make only a single portfolio-allocation choice, then she would take positions in twinned problems that result in the same state-contingent wealth lottery. However, in our experiment each participant does not make a single choice but instead a sequence of choices, with only one of these randomly chosen to be paid out. In this case, we can assume that the decision maker adheres to the Independence Axiom in order to conclude that she chooses the same portfolio across twinned problems. This limits the scope of the result to expected-utility preferences, be they risk-averse, risk-loving, or neither.

Yet Tversky and Kahneman (1981) and Rabin and Weizsäcker (2009), among others, demonstrate that participants in experiments often do not make choices on individual problems that take into account the entire set of problems they have to solve—even when explicitly told to do so—but instead make each choice in isolation, neglecting all remaining problems. We capture this idea with a definition similar to that in Rabin and Weizsäcker (2009).

Definition 1. *The decision maker brackets narrowly if she makes each portfolio-allocation choice as if it were her only one: she applies her preference relation to the choice set given by $\mathcal{W}^n = \mathcal{W}(X_1^n, X_2^n; \mathcal{C}^n)$.*

A participant who “brackets narrowly” in this sense chooses the same portfolio across twinned problems to conform to Observation 1 regardless of whether she obeys the Independence Axiom.

Assumption B. *The decision maker brackets narrowly or makes choices to maximise preferences that satisfy the Independence Axiom.*

Together Assumptions A and B imply Observation 1 in our experiment where only a single choice is actually paid out to subjects. These assumptions will be enough for the main results presented below. But for further analyses, we make further restrictions.

A standard assumption about preferences under uncertainty is that people dislike risk.

Assumption C. *The decision maker is risk averse.*¹²

¹²Formally, for every lottery L , the decision maker weakly prefers the degenerate lottery paying $E[L]$ with certainty to L .

Because each asset in our experiment pays out a positive amount in each state of the world, a subject who maximised Kahneman and Tversky (1979) loss-averse preferences over lab winnings—using a zero reference point—would satisfy Assumptions A-C.

In three out of four pairs of twinned problems (Choices 3-8), Assumptions A-C suffice to make unique predictions about subjects' choices. For instance, consider Choice 3 and Choice 4. In both of them, any portfolio leads to a payoff of 12 in state 1 and 24 in state 4. In Choice 3, investing α_E in Asset E and $1 - \alpha_E$ in Asset F gives payoffs $24\alpha_E + 12(1 - \alpha_E) = 12 + 12\alpha_E$ in state 2 and $12\alpha_E + 24(1 - \alpha_E) = 24 - 12\alpha_E$ in state 3. Since states 2 and 3 occur with the same probability, any risk averter prefers to receive the expected value of her money payoff across the two states, 18, in each state to any available lottery. This can be achieved by choosing $\alpha_E = \frac{1}{2}$. Each of Choices 3 through 8 has the feature that in two states the decision maker can do nothing to hedge her risk while in the remaining two she can perfectly hedge her risk just as in this case.

In Choices 1 and 2, risk aversion alone does not suffice to identify subjects' optimal choice. In this case, adding Assumption D suffices to make a unique prediction.

Assumption D. *The decision maker is arbitrarily close to risk neutral.*¹³

Observation 2. *Under Assumptions A-D, for each n , $m_{\succeq}(\mathcal{W})$ minimises the variance in portfolio earnings in each choice \mathcal{W}^n .*

? contains a more general theorem along these lines and its proof.¹⁴ Intuitively, uniform arbitrary closeness to risk neutrality implies arbitrary closeness to constant-absolute risk aversion (CARA) preferences, which collapse to lexicographic mean-variance preferences as the decision maker approaches risk neutrality.

¹³Formally, consider a sequence of expected utility maximisers with C^2 , concave Bernoulli utility functions $(u_n)_{n \in \mathbb{N}}$ and their associated Arrow-Pratt coefficient-of-absolute-risk aversion functions $(r_n)_{n \in \mathbb{N}}$. Let α_n denote the portfolio that maximises the expectation of u_n . When $(r_n)_{n \in \mathbb{N}}$ converges uniformly to zero, and the decision maker chooses a portfolio that belongs to the limit of $(\alpha_n)_{n \in \mathbb{N}}$, we say that the decision maker is arbitrarily close to risk neutral.

¹⁴Formally, $(\alpha_n)_{n \in \mathbb{N}}$ converges to the portfolio that maximises lexicographic preferences in the moments, with alternating sign. In the context of our experiment where assets have equal means and there is at most one portfolio with any given variance, this coincides with minimising portfolio variance.

An alternative hypothesis to rationality is that our decision maker neglects the correlation in assets' returns, treating each asset as independent. This violates Assumption A because such a decision maker does not maximise rational preferences over state-contingent payoffs. We model a “hexed” decision maker as one who maximises rational preferences over a transformed, fictitious asset space with two assets $\widehat{X}_1, \widehat{X}_2$ that are uncorrelated. To construct the fictitious asset space, let

$$\widehat{\Omega} = \{1.1, 1.2, 1.3, 1.4, 2.1, 2.2, 2.3, 2.4, 3.1, 3.2, 3.3, 3.4, 4.1, 4.2, 4.3, 4.4\}$$

and suppose that in state a, b , asset \widehat{X}_1 pays $X_1(a)$ and \widehat{X}_2 pays $X_2(b)$, and states in $\widehat{\Omega}$ occur equi-probably. Whereas $X_1, X_2 : \Omega \rightarrow \mathbb{R}$, these new assets $\widehat{X}_1, \widehat{X}_2 : \widehat{\Omega} \rightarrow \mathbb{R}$. Most importantly, whatever the correlation between X_1 and X_2 , \widehat{X}_1 is independent of \widehat{X}_2 by construction: whatever the payout of \widehat{X}_1 , \widehat{X}_2 pays out $X_2(1), X_2(2), X_2(3)$ and $X_2(4)$ with equal probability.

Let $\widehat{\mathcal{W}}^n = \widehat{\mathcal{W}}(\widehat{X}_1^n, \widehat{X}_2^n; \mathcal{C}^n)$ be the set of wealth lotteries feasible for the investor when the assets are $\widehat{X}_1^n, \widehat{X}_2^n$ given the constraints \mathcal{C}^n . Let $\widehat{\mathcal{W}} = \cup_n \widehat{\mathcal{W}}^n$ and $\Delta\widehat{\mathcal{W}}$ be the set of probability distributions over $\widehat{\mathcal{W}}$.

Assumption E. *The decision maker makes choices that maximise a preference relation \succeq over $\Delta\widehat{\mathcal{W}}$ that is complete and transitive. Moreover, for each $n, m_{\succeq}(\widehat{\mathcal{W}}^n)$ is a singleton set.*

When X_1 and X_2 are uncorrelated, \mathcal{W} is equivalent to $\widehat{\mathcal{W}}$, which implies that rational and hexed investors make the same choices.

Observation 3. *When X_1 is independent of X_2 , Assumptions A and B produce the same choices as do Assumptions B and E.*

When assets are correlated, Assumptions B,C,D, and E suffice to make a unique prediction for exactly the reasons that Assumptions A, B,C, and D do above. In particular, in each choice the investor allocates her portfolio to minimise her perceived portfolio variance.

When X_1 is positively correlated with X_2 , a hexed investor underestimates the variance in her portfolio, as shown in the following proposition. Let σ_{12} be the covariance between X_1 and X_2 .

Proposition 1. *When $\sigma_{12} > 0$, Assumptions B,C,D, and E give a (weakly) more equal split of*

portfolio across assets than Assumptions A,B,C, and D. When $\sigma_{12} < 0$, Assumptions B,C,D, and E give a (weakly) more unequal split of portfolio across assets than Assumptions A,B,C, and D.

When assets are positively correlated, a hexed investor overestimates diversification benefit of moving his portfolio from a low-variance asset to a high variance one, which leads her to take a more highly diversified portfolio than a rational decision maker, a form of *false diversification effect*. When assets are negatively correlated, a hexed investor underestimates the diversification benefit of moving her portfolio from a low-variance asset to a high variance one, which leads her to take a less diversified portfolio than a rational decision maker, a form of *hedging neglect*.

We now consider our third behavioral theory, the $1/N$ heuristic. Benartzi and Thaler (2001) have suggested that investors facing a menu of N different mutual funds often use the simple heuristic of investing the fraction $1/N$ of her portfolio in each fund. In the context of our experiment, investing half the portfolio in each of the two choices would lead an investor to make different portfolios in paired problems.

We close this section by introducing a parametric model that combines the three behavioral models. As discussed above, A-D as well as B-E imply lexicographic mean-variance preferences, albeit with different perceived covariance matrices. In settings where risk averters are close enough to risk neutral to maximise mean-variance preferences, we propose a simple parametric alternative that incorporates diminishing sensitivity to variance of any kind by modifying the perceived covariance matrix. While much research has demonstrated people’s aversion to risk—even at its smallest scale—that does not imply that people’s dislike for risk increases linearly, as the Independence Axiom implies in this context. Like Fechner and Weber famously demonstrated with a broad range of physical stimuli, people’s perception of or distaste for risk may increase slower than linearly with variance. We capture this possibility by positing that people maximise mean-variance preferences using the following transformed covariance matrix:

$$V(k, l) = \begin{pmatrix} (\sigma_1^2)^l & k \cdot \text{sgn}(\sigma_{12})|\sigma_{12}|^l \\ k \cdot \text{sgn}(\sigma_{12})|\sigma_{12}|^l & (\sigma_2^2)^l \end{pmatrix}.$$

The parameter $k \in [0, 1]$ appears in the off-diagonal elements and captures people’s responsiveness

to covariance. The lower is k , the less people attend to covariance. The parameter $l \geq 0$ captures how people’s dislike for risk depends upon its level. The lower is l , the more people’s marginal dislike for risk diminishes with its level.

This transformed covariance matrix allows us to nest our several hypotheses about behavior.

RAT: Subjects satisfy Assumptions A-D, i.e., $k = l = 1$.

HEX: Subjects satisfy Assumptions B-E, i.e., $k = 0$ and $l = 1$.

1/N: Subjects use the 1/N heuristic, i.e., $l = 0$.¹⁵

Finally, when $k = 0$ and $l < 1$, behaviour conforms to a hexed investor who also exhibits diminishing sensitivity to variance. We return to this parametric model in Section 5 but first turn to the main experimental result.

4 Raw Data and Analysis of Means

This section shows the raw data of the experiment and tests whether Assumptions A and B are satisfied. Overall, the data provides evidence consistent with both directions of correlation neglect (HEX and 1/N).

Figures 1 to 4 depict all choices by the participants. Each figure has the 146 participants’ choices in two paired problems, $\{n - 1, n\}$ for n even, indicated by the small markers. (The large markers are the predictions of the different behavioral assumptions, explained below.) The horizontal axis in each figure measures investment in the first asset, α_1^n , in the even-numbered choice problem—the one with correlated assets. The vertical axis measures investment in the first asset in the twinned odd-numbered problem expressed in the action space of problem n : $\hat{\alpha}_1^{n-1}$ is the share that must be invested in problem n to achieve the same distribution of earnings as α_1^{n-1} from problem $n - 1$. By construction of the twinned-problem design, this share is uniquely determined for each feasible

¹⁵Since $V(k, 0) = \begin{pmatrix} 1 & k \\ k & 1 \end{pmatrix}$, someone with $l = 0$ perceives no difference between the two assets and therefore invests half of her portfolio in each.

investment in problem $n - 1$. Formally, for each $\alpha_1^{n-1} \in \mathcal{C}^{n-1}$ there exists a unique $\hat{\alpha}_1^{n-1} \in \mathcal{C}^n$ such that

$$\forall x \in \mathbb{R}, \forall \omega \in \Omega, \Pr [W^{n-1}(\omega) = x | \alpha_1^{n-1}] = \Pr [W^n(\omega) = x | \hat{\alpha}_1^{n-1}]$$

and thus each portfolio that the investor might choose in choice problem $n - 1$ corresponds to a unique portfolio in problem n . By Observation 1, the null hypothesis that participants correctly perceive covariance gives a simple prediction for the figure: under Assumptions A and B (which implicitly prescribe that correlation is fully understood), the two shares $\hat{\alpha}_1^{n-1}$ and α_1^n are identical, meaning that the data lie on the 45-degree line. The data are also summarized in terms of mean and standard deviation in Table 2, second column.¹⁶




Figure 1: Distribution of investments in Choices 1 and 2. Horizontal axis: α_1^2 . Vertical axis: $\hat{\alpha}_1^1$.

Large {red, green, blue, x} markers represent {RAT, HEX, 1/N, parametric} predictions.

¹⁶For Choice 6, the prediction of the 1/N model violates the a lower bound of 40 points for the first-listed asset. In the figures, in Table 2 and all of the subsequent analysis, we therefore set the 1/N prediction to 40 points.




Figure 2: Distribution of investments in Choices 3 and 4. Horizontal axis: α_1^4 . Vertical axis: $\hat{\alpha}_1^3$. Large {red, green, blue, x} markers represent {RAT, HEX, 1/N, parametric} predictions.




Figure 3: Distribution of investments in Choices 5 and 6. Horizontal axis: α_1^6 . Vertical axis: $\hat{\alpha}_1^5$. Large {red, green, blue, x} markers represent {RAT, HEX, 1/N, parametric} predictions.




Figure 4: Distribution of investments in Choices 7 and 8. Horizontal axis: α_1^8 . Vertical axis: $\hat{\alpha}_1^7$.

Large {red, green, blue, x} markers represent {RAT, HEX, 1/N, parametric} predictions.

n	Data mean (std. dev.)	RAT ^{n}	HEX ^{n}	(1/N) ^{n}
1	0.254(.113)	0.4	0.4	0.25
2	0.381(.138)	0.4	0.333	0.5
3	0.126(.306)	0	0	0
4	0.444(.234)	0	0.333	0.5
5	0.839(.053)	0.833	0.833	0.833
6	0.778(.127)	0.833	0.929	0.667
7	0.328(.225)	0.25	0.25	0.25
8	0.545(.267)	0.25	0.357	0.5

Table 2: Proportions of investment in first-listed asset.

Data and benchmark model predictions, separated
by choice problem

Figures 1-4 show strong and systematic deviations from Assumptions A and B. Not only do only very few observations lie on the 45-degree lines, but there are also patterns in the data that

cannot be driven by unsystematic disturbances. Corresponding statistical tests soundly reject the prediction of Assumptions A and B that behavior is unchanged within pairs of twinned problems. For each of the four pairs, Wilcoxon matched-pairs sign-rank tests reject the hypothesis of identical distributions of portfolios at a significance level of 0.001. This shows that the use of correlated versus uncorrelated asset returns has a significant impact on choices.

As another measure of the accuracy of the null hypothesis of correct perception of correlation we ask how many participants make portfolio choices that are *exactly* identical between twinned problems. The answer is contained in Table 3 showing that the large majority of the participants never or almost never choose portfolios that are identical between twinned problems. 97.2% of the participants choose a different set of portfolios between two twinned problems weakly more often than they choose the same portfolio. 60.9% never choose the same portfolio twice.

#	Freq.	%	Cum. %
0	89	60.9	60.9
1	37	25.3	86.2
2	16	11.0	97.2
3	3	2.1	99.3
4	1	0.7	100.0
Total	146	100.0	

Table 3: Frequencies of choosing identical portfolios in twinned choice problems (out of 4)

Of course, the deviations may be influenced by different sources and we must to be careful not to interpret random deviations from the 45-degree line as evidence of correlation neglect. We therefore turn to statistical estimations that allow us to conclude that the deviations from the prediction of Assumptions A and B are indeed systematic in the way that we hypothesize. The first such estimation is an OLS regression that summarizes the statistical connection between the data and extreme predictions HEX and $(1/N)$. These predictions are indicated in the figures—the green and blue marker, respectively—as well as in columns 4 and 5 of Table 2. As a benchmark prediction, the figures and Table 2 also contain the “rational” prediction RAT (red marker) of a

risk averter who understands the correlation structure (Assumption C).¹⁷ The dependent variable in the regression is the vertical distance of the data points from the 45-degree line in the figure. The explanatory variables are, analogously, the vertical distances of the two predictions from the 45-degree line, $\widehat{\text{HEX}}^{n-1} - \text{HEX}^n$ and $\widehat{1/N}^{n-1} - (1/N)^n$, respectively.¹⁸

	$\widehat{\alpha} - \alpha$		
	(1)	(2)	(3)
$\widehat{\text{HEX}}^{n-1} - \text{HEX}^n$	0.85 (.08) ^{***}	–	0.18 (.08) ^{***}
$\widehat{1/N}^{n-1} - (1/N)^n$	–	0.63 (.05) ^{***}	0.56 (.04) ^{***}
R^2	0.21	0.35	0.36
# of obs.	584	584	584

Table 4: OLS regression of deviation from 45-degree line on the predicted deviation by two models HEX and 1/N. Standard deviations in parentheses, clustered by subject.

The table shows that both extreme models of correlation neglect are predictive of the data means, as their coefficients are statistically significant both individually (in univariate regressions) and when controlling for the respective other prediction. The 1/N model has the larger predictive power. The regressions results are also consistent with the data feature that the deviations from the 45-degree line are strong especially when HEX and 1/N move together. This can be seen by inspection of the models' predictions in Table 2: in Choices 4 and 8, the predictions of HEX and 1/N are both on the same side of the benchmark model RAT. In these two problems the data means also deviate from RAT in the same direction. More generally, most participants' deviations in these problems from the 45-degree line are in the same direction as the two stylized models. In contrast, in the two other choice problems, the two stylized models move in opposite directions away from RAT, and the data are also closer to RAT.

¹⁷Where needed to make a unique prediction, the degree of risk aversion is also assumed to be close to zero, as explained in Section 3 (Assumption D).

¹⁸The estimated model is

$$\widehat{\alpha}_1^{n-1,i} - \alpha_1^{n,i} = \beta_2(\widehat{\text{HEX}}^{n-1,i} - \text{HEX}^{n,i}) + \beta_3(\widehat{1/N}^{n-1,i} - (1/N)^{n,i}) + \nu^{n,i}$$

where n is even and i indexes the participant. Since the model is differenced, it excludes a constant term.

To see the partial predictive power of all three models, we now consider only correlated tasks (where the three model predictions differ), and regress α on all three model predictions including RAT.

	α			
	(1)	(2)	(3)	(4)
RAT	-0.44 (.07)***	0.18 (.04)***		-0.13 (.06)**
HEX	1.35 (.06)***	—	0.37 (.04)***	0.53 (.06)***
1/N	—	0.87 (.04)***	0.65 (.05)***	0.61 (.04)***
R_u^2	0.85	0.87	0.88	0.88
# of obs.	584	584	584	584

Table 6: OLS regression of investments in even-numbered choice problems on model predictions. Standard deviations in parentheses, clustered by subject. Due to absence of constant, goodness-of-fit is measured by the uncentered R_u^2 .

The table shows that when only RAT and 1/N are included, RAT has positive explanatory power and a coefficient of the expected sign. But strikingly, RAT has a negative coefficient once HEX is included. This points again at strong misperceptions of the implied correlation. An important caveat is that the above OLS regressions and Table 2 focus on the data means, implicitly viewing all participants as homogeneous—whereas the data show strong patterns of heterogeneity. Section 5 therefore shows model estimates where heterogeneity between the different participants is allowed.

5 Estimates of Heterogeneous-Type Models

In this section we estimate the parameters of the model presented in Section 3, where $l \geq 0$ measures how strongly participants register differences among the (co-)variance terms, while $k \in [0, 1]$ measures the extent to which participants consider covariance at all. It is straightforward to show that a decision maker who perceives the covariance matrix in problem n to be $V(k, l)$ minimises her perceived variance in portfolio earnings by investing the fraction $\alpha_1^n(k, l) = \frac{(\sigma_2^2)^l - k|\sigma_{12}|^l \cdot \text{sgn}(\sigma_{12})}{(\sigma_1^2)^l + (\sigma_2^2)^l - 2k|\sigma_{12}|^l \cdot \text{sgn}(\sigma_{12})}$

of wealth in the first asset.¹⁹ We begin by estimating separate maximum-likelihood values of (k, l) for each individual separately, using her four decisions that involve correlated assets.²⁰

For the estimation, we assume that participant i when making her choice $\alpha_1^{n,i}$ has her own individual-specific parameter vector (k^i, l^i) and follows the prediction $\alpha_1^{n,i} = \alpha_1^n(k^i, l^i)$ except insofar as she is subject to the disturbance term ϵ_i^n : she chooses the investment level $60 \cdot (\alpha_1^n(k^i, l^i) + \epsilon_i^n)$, rounded to the nearest integer value. We assume that the distribution of ϵ_i^n is logistic and i.i.d. across i and n . The likelihood of observing the participant's quadruple $\alpha_1^i = (\alpha_1^{2,i}, \alpha_1^{4,i}, \alpha_1^{6,i}, \alpha_1^{8,i})$ is therefore given by

$$L(\alpha_1^i | k^i, l^i) = \prod_{n=2,4,6,8} \Pr \left[\alpha_1^n(k^i, l^i) + \epsilon_i^n = \alpha_1^{n,i} \right]$$

Maximizing this likelihood separately for each individual gives the distribution of estimates depicted in Figure 5.




Figure 5: Estimates of subject-specific parameters (k, l) , for all 146 participants.

¹⁹The proof of Proposition 1 in Appendix C contains this derivation for the case where $l = 1$; the extension to general l is immediate.

²⁰All estimations in this section include only Choices 2, 4, 6 and 8 because in most odd-numbered tasks the predictions coincide for all values of k and l .

The figure suggests the existence of different perceptions of the covariance structure: some participants behave as if they completely ignore all variances and covariance as they have an estimated parameter $l^i = 0$. Others have “more rational” estimated value of l^i alongside an estimated k^i of zero, indicating that they ignore off-diagonal elements of the covariance matrix.

To summarize the data patterns more concisely—and enable statistical inference—we reduce the number of free parameters by classifying participants into T different types, where each type $t \in \{1, \dots, T\}$ has a different parameter vector (k^t, l^t) . We estimate the proportions of the T types $\{\pi^t\}_{t=1\dots T}$ (restricted to sum to one) together with their behavioral parameters $\{(k^t, l^t)\}_{t=1\dots T}$. The model is thus a mixture model where the likelihood of observing the collection of participants’ choice vectors $\alpha_1 = \{\alpha_1^i\}_{i=1}^{146}$ is:

$$L(\alpha_1 | \{k^t, l^t, \pi^t\}_{t=1\dots T}) = \prod_{i=1}^{146} \sum_{t=1}^T \pi^t \prod_{n=2,4,6,8} \Pr(\alpha_1^n(k^t, l^t) + \epsilon_i^n = \alpha_1^{n,i})$$

The model generates different results depending on the number of types T included. We ran the estimations for up to $T = 5$ and report the results for $T \in \{1, 2, 3\}$ in the following Table 7. For $T = 4, 5$, the results are qualitatively similar to the case $T = 3$ but the likelihood is scarcely improved relative to $T = 3$.

T	ll^*	(l^1, k^1, π^1)	(l^2, k^2, π^2)	(l^3, k^3, π^3)
1	-2296.6	(0.43, 0, 1)	–	–
2	-2288.4	(0.39, 0, 0.91)	(1.1, 1, 0.09)	–
3	-2286.2	(0.45, 0, 0.58)	(1.09, 1, 0.1)	(0, –, 0.3)

Table 7: Estimates of parametric mixture model with $T \in \{1, 2, 3\}$ types.

The table’s first row of entries shows the result under the restriction that all participants belong to a single type, $T = 1$. The best-fitting indicate that the estimated single type in this model completely ignores the off-diagonal elements of the covariance matrix ($k^1 = 0$) and only partially reacts to differences in the assets’ variances ($l^1 = 0.43$). This resembles the prediction HEX yet echoes 1/N through substantial insensitivity to differences in variance. This model is rejected in favor of the two-type model $T = 2$ ($p < 0.01$, likelihood-ratio test), where the best-fitting two co-existent types are quite different in nature: the primary type ($\pi^1 = 0.91$) resembles

the one estimated in the single-type model, while the secondary type bears much closer semblance to the rational prediction of full appreciation of correlation. Finally, with $T = 3$, the best-fitting parameter constellation also includes a pure $1/N$ type. The result of the $T = 3$ estimation is thus not too far from a model consisting of the three archetypical types HEX, RAT and $1/N$ all appear: each of these three extreme types is reasonably well approximated by one of the three estimated types.²¹ However, the $T = 2$ specification is not rejected in favor of $T = 3$ ($p=0.221$): adding a $1/N$ type does not significantly improve fit over the rational and (modified) Hex type.

In sum, our most reliably estimated single type is that estimated when $T = 1$. We depict that prediction in Figures 1 to 4 with an “x”. In each case, the x lies near the centre of the data point cloud (and closely matches the sample means presented in Table 2).

6 Additional Demonstrations of Correlation Neglect

We designed two additional experimental tasks to elicit correlation neglect in ways other than those of the previous sections. Different participants receive different instructions for these additional tasks. As noted earlier, because all of the additional tasks appear after the main part of the experiment, these differences in instructions could not affect the results discussed in the previous sections.

A notable modification from the previously discussed problems is that the expected return varies across assets. Yet all additional tasks are portfolio-choice problems of the same general format as in the experiment’s main part. Only very brief additional instructions are therefore needed between the two parts of the experiment. Section 6.1 describes the first set of additional tasks, presented to three different subgroups of participants after the main experiment. Section 6.2 describes the second set of tasks, presented again to other subgroups of participants. For the task described in

²¹Estimating a model that includes only the three extreme types, HEX, RAT and $1/N$, in unknown proportion yields estimates markedly different from those in Table 7 and reminiscent of the linear regression in Table 4: the $1/N$ type gets weight $\pi^{sc1/n} = 0.75$, while the HEX type gets only weight $\pi^{scHex} = 0.16$. This shows that the first-listed type estimate in Table 7, which is similar to HEX, would be much less accurate if $l = 1$ was assumed instead of leaving l unrestricted.

Section 6.2, a set of new participants were used, as described later.

6.1 Falsely Perceived Stochastic Dominance

In this example, we construct two portfolios, one of which appears to first-order-stochastically dominate the second to someone who neglects correlation. Unlike in the previous (nearly) continuous-choice problems, the decision-maker here chooses discretely between two fixed portfolios, each one holding a strictly positive amount of both assets. The asset distributions and contents of the two available portfolios have the property that a decision-maker who ignores covariance, like the hexed agent of Section 3 (Assumptions B-E), falsely perceives a first-order stochastic dominance relation between the two available portfolios.

	[1]	[2]	[3]
U	30	20	12
V	10	12	30

Please indicate your preferred investment, by ticking the box:

- I invest 52 points in Asset U and 8 points in Asset V .
- I invest 26 points in Asset U and 34 points in Asset V .

Clearly, U has a higher return distribution than V , with a difference in means of $\frac{10}{3}$. This makes the first choice relatively more attractive due to its high weight on U . On the other hand, the negative correlation between the assets may entice a sufficiently risk averse decision-maker to take the safer second choice.

Crucially, the task is designed such that a hexed decision-maker who ignores the negative correlation would never choose the safe option. It is straightforward to show that if the two assets are perceived to be independent—in the notation of Section 3, if the preferences relation \succeq is defined over $\Delta\widehat{\mathcal{W}}$ (Assumption E) and the decision-maker considers assets \widehat{U}, \widehat{V} instead of U, V —then the first choice first-order stochastically dominates the second choice.

The task was presented to 30 subjects, 29 of whom (97%) chose the first option. This result is consistent with the presence of hexed decision-makers. However, it may also be driven by rational preference maximization—perhaps all of these 29 participants correctly perceived the covariance structure but are not sufficiently risk averse to avoid the relatively higher risk. To demonstrate the effect of the perception of the covariance structure, we therefore repeat the task with 39 different participants, presenting the same portfolios in a way that does away with any need to contemplate covariance.

	[1]	[2]	[3]
U'	27.3	18.9	14.4
V'	15.5	18.7	22.2

Please indicate your preferred investment by ticking the box:

- I invest 60 points in Asset U' and 0 points in Asset V' .
- I invest 0 points in Asset U' and 60 points in Asset V' .

The reader can verify that the return distributions of assets U' and V' are identical to the distributions resulting from the two choice options in the first problem, involving assets U and V (up to rounding error). Therefore the two choice options in the second problem, each allowing only investments of all wealth in one asset, are identical to those in the first problem—a pure framing variation. Without any diversification opportunity, covariance between the available options is irrelevant to the second problem, so that even hexed decision-makers would not err. The pair of problems therefore produces the possibility of preference reversals for hexed agents. Indeed we see a significant change in behavior: in the second framing variation of the problem, only 20 out of 39 participants (51%) choose the first option.²² The difference between results in the two framing variations is significant at $p < 0.01$ (Wilcoxon two-sample test).

²²The result is further corroborated by additional framing variations that replace only one of the two portfolios consisting of Assets U and V by a portfolio whose state-contingent return is explicitly presented. 34 subjects (all of whom are included in the analysis of Sections 2-5 but not that of Section 6 up until this point) face the same binary task twice, in two different framing variations hard to recognize as identical. In the first variation, the first option

6.2 Ignorance of Arbitrage

Not only may people who neglect correlation perceive first-order-stochastic dominance where it is not, but they may fail to perceive it where it is. In our portfolio-choice setting, they may choose one portfolio over a second that pays higher returns in every possible state. We explore this possibility through two portfolio-choice problems that each use three assets. Each problem is designed such that there exist portfolios combining two assets that statewise dominate the third. Any investment in the third therefore violates “arbitrage freeness”. These two tasks are presented at the end of the experiment (like those in Section 6.1) to 40 participants of the main experiment described in Section 2 and to 96 new participants. The 96 new participants first face the same instructions and tasks as in the main experiment, the only difference being that in the four problems involving correlated assets (Choices 2, 4, 6, 8) their choice options are restricted to two portfolios. (The results of these restricted tasks essentially confirm the results of the previous sections and we skip them for brevity.) A total of 136 participants completed the two tasks described here.

11. Invest each of your 60 points in Asset U'' , Asset V'' or Asset W'' , as given below.

	[1]	[2]	[3]
U''	15	38	7
V''	39	5	16
W''	20	15	13

In the first of the two additional tasks, Choice 11, participants choose their portfolio weights on Assets U'' , V'' and W'' , as depicted above.²³ One can check easily that asset W'' is dominated by a continuum of combinations of U'' and V'' —in particular it is dominated by $\frac{1}{3}U'' + \frac{2}{3}V''$ (with 52 points in U and 8 points in V) appears as the option of investing everything in U' . In the second variation, the second option (26 points in U and 34 points in V) is replaced by the option of investing everything in V' . Under both variations, correlation neglect is irrelevant, so that a hexed agent should behave exactly as in the case where both choice options are shown as U' and V' . Indeed, 18 out of 33 participants (55%) opt for the first choice option under the first framing variation (one participant did not fill in the decision sheet) and 16 out of 34 participants under the second framing variation.

²³In the experiment the asset labels are U , V and W (symbols that we already used for other participants' problems).

a expected return difference of 4 Euros—rather a large difference). This arbitrage derives from the negative correlation of U'' and V'' that allows an effective insurance against their respective low-return outcomes. But many participants appear to neglect the hedging opportunity: the average portfolio weight of W'' is 0.257, close to the uniform weight. (The average weights of U'' and V'' are 0.346 and 0.397 respectively.) Moreover, the distribution of investments reveals that 84 out of 136 participants (75%) invested a strictly positive wealth share in W'' . That is, a large majority of subjects makes a dominated choice in this task.

The last-mentioned result is even more pronounced in the second additional task:

12. Invest each of your 60 points in Asset X, Asset Y or Asset Z, as given below.

	[1]	[2]	[3]
X	14	22	12
Y	27	5	16
Z	18	16	13

Here, Asset Z is dominated by a continuum of combinations of X and Y —for example by $\frac{2}{3}X + \frac{1}{3}Y$. Only 19 out of 136 participants (14%) invested nothing in Z and thereby avoided the dominated investment. The remaining 86% of subjects appear not to appreciate the dominance relation that arises due to negative correlation of X and Y . The average portfolio weights of X , Y and Z are 0.399, 0.308 and 0.293 respectively.

7 Conclusion

This paper has presented a sequence of portfolio-choice experiments demonstrating that people’s asset allocations depend strongly upon how financial assets are framed, in particular whether the same asset span is generated using correlated or uncorrelated assets. At the broadest level, participants systematically violate the “reduction of compound lotteries” property that underlies the theory of choice under uncertainty: the same portfolio appeals to people differently depending upon how it is constructed.

Our experiments also provide strong evidence that people succumb to a specific form of framing effect in financial decision-making, namely they neglect correlation in asset returns. In part, this may derive from the form of $1/N$ rule that Benartzi and Thaler (2001) observe in financial decision-making, which also has been uncovered in other contexts. Simonson (1990) as well as Read and Loewenstein (1995) have suggested that people use strategies of “naive diversification” by instinctively diversifying their choices despite having underlying preferences that do not warrant such diversification. For instance, a trick-or-treater who prefers Mars to Snickers bars may too frequently choose Snickers over Mars so as not to end up with an unbalanced candy supply. Indeed, restaurant-goers often seem reluctant to order dishes that previously ordered by their companies, even when they know that they will not share dishes.

In settings where people’s choice sets are imperfectly observed, the $1/N$ heuristic provides the modeler with little guidance as to what predictions to make about behaviour. By contrast, our model of correlation neglect (the “hexed” agent) predicts how people’s choices depart systematically from what they would choose had they taken correlation into account fully. In a two-asset setting, Proposition 1 establishes that when correlation is positive, hexed investors diversify more than standard theory predicts—exaggerating the benefit of diversification—while when correlation is negative they diversify too little—under-appreciating its benefit.

People who systematically under-appreciate correlation have a lower willingness to pay for assets that hedge their portfolio risk and a higher willingness to pay for assets that magnify their portfolio risk. Without perfect arbitrage (see, e.g., Shleifer (2000) for limits to arbitrage), a systematic under-appreciation of correlation may thus cause hedging opportunities to be systematically underpriced in financial transactions or markets.

Finally, our data analysis provides novel evidence that once endowed with risk, people’s dislike for more risk decreases. Kőszegi and Rabin (2007) propose a theory of reference-dependent risk preferences under which people with deterministic reference lotteries are more averse to risk than those with stochastic reference lotteries. In the context of our experiment, participants who bracket narrowly may become inoculated by unavoidable risk against the pain of taking on discretionary risk. Experiments designed explicitly to test for diminishing sensitivity to risk (variance) might

help researchers better understand the scale and scope of this phenomenon.

References

- Benartzi, S.: 2001, Excessive extrapolation and the allocation of 401(k) accounts to company stock?, *Journal of Finance* **56**(5), 1747–1764.
- Benartzi, S. and Thaler, R.: 2001, Naive diversification strategies in defined contribution saving plans, *American Economic Review* **91**(1), 79–98.
- Brunnermeier, M.: 2009, Deciphering the liquidity and credit crunch 2007-2008, *Journal of Economic Perspectives* **23**(1), 77–100.
- Eyster, E.: 2010, Near risk neutrality and lexicographic preferences in the moments, *mimeo* .
- Gubaydullina, Z. and Spiwoks, M.: 2009, Portfolio diversification: an experimental study.
- Hellwig, M.: 2009, Systemic risk in the financial sector: An analysis of the subprime-mortgage financial crisis, *De Economist* **157**(2), 129–207.
- Kahneman, D. and Tversky, A.: 1979, Prospect theory: An analysis of decision under risk, *Econometrica* **47**(2), 263–292.
- Kallir, I. and Sonsino, D.: 2009, The neglect of correlation in allocation decisions, *Southern Economic Journal* **75**(4), 1045–1045.
- Kőszegi, B. and Rabin, M.: 2007, Reference-dependent risk attitudes, *American Economic Review* **97**(4), 1047–1073.
- Kroll, Y., Levy, H. and Rapoport, A.: 1988, Experimental tests of the separation theorem and the capital asset pricing model, *American Economic Review* **78**(3), 500–519.
- Polkovnichenko, V.: 2005, Household portfolio diversification: A case for rank-dependent preferences, *Review of Financial Studies* **18**(4), 1467–1502.

- Rabin, M. and Weizsäcker, G.: 2009, Narrow bracketing and dominated choices, *American Economic Review* **99**(4), 1508–1543.
- Read, D. and Loewenstein, G.: 1995, Diversification bias: Explaining the discrepancy in variety seeking between combined and separated choices, *Journal of Experimental Psychology: Applied* **1**(1), 34–49.
- Shleifer, A.: 2000, *Inefficient markets: an introduction to behavioral finance*, Oxford University Press.
- Simonson, I.: 1990, The effect of purchase quantity and timing on variety seeking behavior, *Journal of Marketing Research* **32**, 150–162.
- Tversky, A. and Kahneman, D.: 1981, The framing of decisions and the psychology of choice, *Science* **211**(4481), 453–458.

8 Appendix A: Instructions

Welcome!

You are about to participate in an experiment in decision making. Universities and research agencies have provided the funds for this experiment.

In this experiment we will first ask you to read instructions that explain the decision scenarios you will be faced with. We will also ask you to answer questions that test your understanding of what you read. Finally, you will be asked to make decisions that will allow you to earn money. Your monetary earnings will be determined by your decisions and by chance. All that you earn is yours to keep and will be paid to you in private, in cash, after today's session.

Only for coming here and completing the experiment, you will also receive a fixed participation fee of EUR 5.00. The earnings that you make during the experiment will be added to this amount.

It is important to us that you remain silent and do not look at other people's work. If you have any questions or need assistance of any kind, please raise your hand, and an experimenter will come

to you. If you talk, exclaim out loud, etc., you will be asked to leave and will forfeit your earnings. Thank you.

The determination of your pay-out

We ask you to make 11 decisions in total during this experiment. Of these 11 decisions, only one will be actually paid out. This will be determined by a random draw at the end of today's session, when you will make a draw from a stack of 11 cards. The card that you draw will determine which decision will be paid out to you. Every participant in the room will make a separate draw, so different decisions will be paid out to different participants.

Description of the decisions

All 11 decisions follow essentially the same format. In each of them, you will allocate 60 points among two different assets. An asset is an investment opportunity that yields a return. The asset returns are probabilistic, meaning that at the time you invest you cannot be sure how much return an asset yields. Each asset has at most 4 different return values. For example, a given asset may have the following possible return values, in Euros:

Asset A 10 28 12 23

If you invest all 60 points in this asset, you would receive the full return that is indicated by the return values. In this example, if you invest all points in Asset A, you would either receive €10 or €28 or €12 or €23. You would not know which of the 4 possible returns you would get.

In each decision, there are two available assets that have different distributions of the probabilistic return, and you can allocate your 60 points between the two assets in the way that suits you best. If you invest less than the full 60 points in a given asset, you receive proportionally less than the full return. For example, if you invest 30 points in Asset A, you would receive either €5 or €14 or €6 or €11.50 from Asset A, plus another probabilistic amount from investing the remaining 30 points in the second asset.

For both assets, the actual return depends on the value of a single random outcome. There are 4 possible values of this outcome, and the return of each asset is different for the 4 different outcome values. The four possible values of the outcome all have the same probability of 1/4. We implement this by flipping 2 coins. One of our two coins has labels # and & on its two sides.

The other coin has two labels $*$ and \wedge . When we flip both of the two coins, there are 4 possible combinations of labels that can come up:

$$\#^* \quad \#\wedge \quad \&^* \quad \&\wedge$$

Each of these four outcomes has the exact same probability of $1/4$ (or 25%). For both of the assets, the outcome depends on which combination of labels comes up. We will toss the two coins only once at the end of the experiment, and this double coin toss will determine the return of both assets in each of the 11 decisions.

For example, one decision that you may face could look like this:

Invest each of your 60 points in either Asset A or Asset B, as given below.

	$\#^*$	$\#\wedge$	$\&^*$	$\&\wedge$
<i>Asset A</i>	15	30	15	20
	$\#^*$	$\#\wedge$	$\&^*$	$\&\wedge$
<i>Asset B</i>	10	20	30	10

In this example, how much you earn depends on how many points you invest in Asset A and Asset B and on the outcome of the two coin tosses. For instance, if you invest 20 points in Asset A and 40 points in Asset B – i.e. you invest one-third of your points in Asset A and two-thirds in Asset B – and if the value of the random outcome is $\&^*$, then your overall earnings would be:

$$(1/3) \times EUR\ 15 + (2/3) \times EUR\ 30 = EUR\ 25$$

Please note that you would be paid this amount only if this particular decision is paid out. In the actual experiment, only one of the 11 decisions will be paid out, as described above.

Also, importantly, note again that only one random outcome is drawn (here, $\&^*$), which determines the return of both assets.

You can easily calculate other possible earnings, for different combinations of token allocations and random outcomes. For example, if you invest 48 points in Asset A and 12 in Asset B—hence, four-fifths in Asset A—and if the random outcome value is $\#\wedge$, your return is:

$$(4/5) \times EUR\ 30 + (1/5) \times EUR\ 20 = EUR\ 28$$

Or, if you choose the same division (48/12 in Assets A/B) but the random outcome value is #*, your return is:

$$(4/5) \times EUR\ 15 + (1/5) \times EUR\ 10 = EUR\ 14.$$

The 11 decisions differ from one another as well as from the above example in two ways:

Firstly, the probabilistic returns of the assets are different. (But they all have at most 4 possible values.)

Secondly, not all investment combinations are available in each case. For example, some decisions require you to invest at least some minimum amount in a given asset. This requirement appears clearly in the instruction line of each separate decision. For example, the precise wording of the instruction may be as follows.

Invest each of your 60 points in either Asset A or Asset B, as given below.

Note: At least 20 points must be invested in Asset A.

	#*	#^	&*	&^	
<i>Asset A</i>	15	20	15	30	<i>Investment : ...points (≥ 20)</i>
	#*	#^	&*	&^	
<i>Asset B</i>	10	0	30	10	<i>Investment : ...points</i>

Please ensure that your two investments sum to 60 points.

For all decisions, only positive token investments are feasible. That is, at least 0 points must be invested in each asset in any case, even if this requirement is not explicitly mentioned in the respective instruction lines.

Before proceeding to the actual decisions, we will ask you to complete an understanding test of the instructions.

Please wait until the understanding test is distributed. If you have any questions about the instructions up to here, please raise your hand.

Understanding test

Please record your code number on this sheet, as well as on all subsequent sheets during the experiment

Consider the following investment decision and answer the questions (1) to (3) below. You will only be allowed to continue with the experiment after answering correctly. If you have a question of any kind, please raise your hand.

Invest each of your 60 points in either Asset A or Asset B, as given below.

Note: At least 10 points must be invested in Asset A.

	$\#^*$	$\#^{\wedge}$	$\&^*$	$\&^{\wedge}$	
<i>Asset A</i>	30	20	15	10	<i>Investment : ---points (≥ 10)</i>
	$\#^*$	$\#^{\wedge}$	$\&^*$	$\&^{\wedge}$	
<i>Asset B</i>	40	0	5	30	<i>Investment : ---points</i>

Please ensure that your two investments sum to 60 points.

Questions:

(1) Suppose that this decision is paid to you at the end of the experiment and that you invest 60 points in Asset A and 0 points in Asset B. If the random outcome value is $\#^{\wedge}$, what do you earn? EUR.....

(2) Suppose that this decision is paid to you at the end of the experiment and that you invest 15 points in Asset A and 45 points in Asset B. If the random outcome value is $\#^{\wedge}$, what do you earn? EUR.....

(3) Suppose that this decision is paid to you at the end of the experiment and that you invest 30 points in Asset A and 30 points in Asset B. If the random outcome value is $\&^*$, what do you earn? EUR.....

Once you finish the understanding test, please wait for instructions for the decisions. If you have a question, please raise your hand. Please make sure that the code number is recorded on the understanding test.

The 11 decisions

In each of the following decisions, please make sure that your investments exactly add up to 60 points. If they do not add up to 60 in the task that is to be paid out, you will only be paid the

smallest possible pay-out in the respective decision.

For decisions 1-10, please indicate your investments next to the respective assets.

1.

Invest each of your 60 points in either Asset A or Asset B, as given below.

$\#^*$ $\#^{\wedge}$ $\&^*$ $\&^{\wedge}$
Asset A 15 21 15 21 *Investment : ___points*

$\#^*$ $\#^{\wedge}$ $\&^*$ $\&^{\wedge}$
Asset B 12 12 24 24 *Investment : ___points*
Please ensure that your two investments sum to 60 points.

2.

Invest each of your 60 points in either Asset C or Asset D, as given below.

Note: At least 30 points must be invested in Asset D.

$\#^*$ $\#^{\wedge}$ $\&^*$ $\&^{\wedge}$
Asset C 18 30 6 18 *Investment : ___points*
 $\#^*$ $\#^{\wedge}$ $\&^*$ $\&^{\wedge}$
Asset D 12 12 24 24 *Investment : ___points (≥ 30)*
Please ensure that your two investments sum to 60 points.

3.

Invest each of your 60 points in either Asset E or Asset F, as given below.

$\#^*$ $\#^{\wedge}$ $\&^*$ $\&^{\wedge}$
Asset E 12 24 12 24 *Investment : ___points*
 $\#^*$ $\#^{\wedge}$ $\&^*$ $\&^{\wedge}$
Asset F 12 12 24 24 *Investment : ___points*
Please ensure that your two investments sum to 60 points.

4.

Invest each of your 60 points in either Asset G or Asset H, as given below.

	#*	#^	&*	&^	
<i>Asset G</i>	12	24	12	24	<i>Investment : ___points</i>

	#*	#^	&*	&^	
<i>Asset H</i>	12	18	18	24	<i>Investment : ___points</i>

Please ensure that your two investments sum to 60 points.

5.

Invest each of your 60 points in either Asset I or Asset J, as given below.

	#*	#^	&*	&^	
<i>Asset I</i>	14	21	14	21	<i>Investment : ___points</i>

	#*	#^	&*	&^	
<i>Asset J</i>	14	14	21	21	<i>Investment : ___points</i>

Please ensure that your two investments sum to 60 points.

6.

Invest each of your 60 points in either Asset K or Asset L, as given below.

Note: At least 40 points must be invested in Asset K.

	#*	#^	&*	&^	
<i>Asset K</i>	14	21	14	21	<i>Investment : ___points (≥ 40)</i>

	#*	#^	&*	&^	
<i>Asset L</i>	14	0	35	21	<i>Investment : ___points</i>

Please ensure that your two investments sum to 60 points.

7.

Invest each of your 60 points in either Asset M or Asset N, as given below.

	#*	#^	&*	&^	
<i>Asset M</i>	12	30	12	30	<i>Investment : ___points</i>

	#*	#^	&*	&^	
<i>Asset N</i>	12	12	30	30	<i>Investment : ___points</i>

Please ensure that your two investments sum to 60 points.

8.

Invest each of your 60 points in either Asset O or Asset P, as given below.

	#*	#^	&*	&^	
<i>Asset O</i>	12	30	12	30	<i>Investment : ___points</i>

	#*	#^	&*	&^	
<i>Asset P</i>	12	18	24	30	<i>Investment : ___points</i>

Please ensure that your two investments sum to 60 points.

9.

Invest each of your 60 points in either Asset Q or Asset R, as given below.

	#*	#^	&*	&^	
<i>Asset Q</i>	16	15	20	21	<i>Investment : ___points</i>

	#*	#^	&*	&^	
<i>Asset R</i>	14	15	19	18	<i>Investment : ___points</i>

Please ensure that your two investments sum to 60 points.

10.

Invest each of your 60 points in either Asset S or Asset T, as given below.

	#*	#^	&*	&^	
<i>Asset S</i>	18	17	19	18	<i>Investment : ___points</i>

	#*	#^	&*	&^	
<i>Asset T</i>	16	16	17	17	<i>Investment : ___points</i>

Please ensure that your two investments sum to 60 points.

[New set of instructions.]

Comment on decision 11: For this decision, you only have the choice between two different combinations of token investments. The two available options will be indicated below the description of the assets.

A further change is that only 3 random outcomes are possible, not 4. Correspondingly, each asset has only 3 possible return values. If you choose the card with number 11, indicating that

decision 11 is paid out to you, then we simply ignore the fourth possible outcome of the double coin toss, when determining your pay-out. That is, if the double coin toss yields $\&^\wedge$, then we will repeat the double coin toss until we get a result of $\#^*$ or $\#^\wedge$ or $\&^*$. The pay-out will then follow the realized outcome. The three outcomes $\#^*$, $\#^\wedge$ and $\&^*$ are therefore the only relevant outcomes in this decision, and they are still equally likely to appear. Effectively, each of the three outcomes will result with probability $1/3$, or one out of three times.

11.

Invest each of your 60 points in either Asset U or Asset V, as given below. Indicate your investment below the description of the assets, choosing from the two available investments.

	$\#^*$	$\#^\wedge$	$\&^*$
<i>U</i>	30	20	12
	$\#^*$	$\#^\wedge$	$\&^*$
<i>V</i>	10	12	30

Please indicate your preferred investment, by ticking the box:

- I invest 52 points in Asset *U* and 8 points in Asset *V*.
- I invest 26 points in Asset *U* and 34 points in Asset *V*.

Once you finish making the decisions, please wait until the experimenter collects the decision sheets. If you have a question, please raise your hand. Please make sure that the code number is recorded on the first decision sheet.

9 Appendix B: Further result on ignored dominance relation

This appendix presents the two remaining investment tasks, Choices 9 and 10, that are presented together with the eight choice problems discussed in Sections 1-5. Both tasks use two assets, where one asset statewise dominates the other.

- 9. Invest each of your 60 points in either Asset Q or Asset R, as given below.

	[1]	[2]	[3]	[4]
<i>Q</i>	16	15	20	21
<i>R</i>	14	15	19	18

10. Invest each of your 60 points in either Asset S or Asset T, as given below.

	[1]	[2]	[3]	[4]
<i>S</i>	18	17	19	18
<i>T</i>	16	16	17	17

The main difference between the two tasks is that in Choice 10, the minimum return of asset *S* weakly exceeds the maximum return of Asset *T*, whereas the same is not true for *Q* and *R* in Choice 9. Therefore, even the hexed agent (Assumptions B to E) would recognise that any strictly positive portfolio weight on *T* is statewise dominated (in the hypothetical state space of the hexed agent) by the portfolio that invests all wealth in *S*. By contrast, in Choice 9 the hexed agent would not even perceive a first-order stochastic dominance relation and may for some (very risk-averse) preferences invest a positive proportion of wealth in *R*, which has illusory diversification benefit.²⁴

The results indeed show a statistically significant difference between the two choice problems, yet the difference is small. In both, participants choose high portfolio weights on the superior assets: in Choice 9, the average weight on *Q* was 0.85, while in Choice 10, the average weight was 0.88 ($N = 242$).²⁵ The difference is statistically significant at $p = 0.004$ (Wilcoxon test). The number of participants who invest a strictly positive proportion of wealth in the dominated asset is also higher in Choice 9 (38%) than in Choice 10 (31%), confirming that the statewise dominance is obeyed more stringently in the choice problem where all hexed agents agree with the optimality of the dominating choice.

²⁴The two choice problems have equal monetary incentives for a rational risk-neutral agent: the first asset has a mean return of 18 and the second-listed asset a mean return of 16.5.

²⁵The reported results include the 96 additional participants who received some of the other problems with a restricted choice set, see Section 6. The results for the main 146 participants are close to identical but show somewhat lower significance levels.

10 Appendix C: Proof

Proof of Proposition 1 Let

$$V = \begin{pmatrix} \sigma_1^2 & k\sigma_{12} \\ k\sigma_{12} & \sigma_2^2 \end{pmatrix}.$$

be the covariance matrix. The variance of a portfolio investing share α in X_1 is then $var(\alpha X_1 + (1 - \alpha)x_2) = \alpha^2\sigma_1^2 + (1 - \alpha)^2\sigma_2^2 + 2\alpha(1 - \alpha)\sigma_{12}$. At an interior minimum, $\alpha_1 = \frac{\sigma_2^2 - k\sigma_{12}}{\sigma_1^2 + \sigma_2^2 - 2k\sigma_{12}}$, so

$$\begin{aligned} \operatorname{sgn}\left(\frac{\partial\alpha_1}{\partial k}\right) &= \operatorname{sgn}\left(-\sigma_{12}(\sigma_1^2 + \sigma_2^2 - 2k\sigma_{12}) + 2\sigma_{12}(\sigma_2^2 - k\sigma_{12})\right) \\ &= \operatorname{sgn}\left(\sigma_{12}(\sigma_2^2 - \sigma_1^2)\right). \end{aligned}$$

Wlog assume that $\sigma_2^2 - \sigma_1^2 \geq 0$ so that $\alpha_1 \geq \frac{1}{2}$. When $\sigma_{12} > 0$, α_1 is increasing in k , meaning that a hexed investor with $k = 0$ chooses a portfolio closer to 50-50 than a rataional agent. Conversely, when $\sigma_{12} < 0$, α_1 is decreasing in k , meaning that a hexer with $k = 0$ chooses a portfolio further from 50-50 than a rational agent. Q.E.D.