

Stable Randomization^{*}

Marina Agranov[†] Paul J. Healy[‡] Kirby Nielsen[§]

Abstract

We design a laboratory experiment to identify whether a preference for randomization defines a stable type across different choice environments. In games and individual decisions, subjects face twenty simultaneous repetitions of the same choice. Subjects can randomize by making different choices across the repetitions. We find that randomization does define a type that's predictable across domains. A sizable fraction of individuals randomize in all domains, even in questions that offer a stochastically-dominant option. For some mixers, dominated randomization is responsive to intervention. We explore theoretical foundations for mixing, and find that most preference-based models are unable to accommodate our results.

Keywords: Randomization; Probability matching; Convex preferences; Stochastic choice; Contingent reasoning

JEL Classification: D81, C91, D89.

^{*}We thank Yaron Azrieli, Ryan Oprea, Pietro Ortoleva, Collin Raymond, and John Rehbeck for helpful comments and suggestions. This study was approved by the Institutional Review Boards at the Ohio State University and Caltech and jointly funded by both institutions.

[†]California Institute of Technology, Division of the Humanities & Social Sciences; *magranov@hss.caltech.edu*

[‡]The Ohio State University, Department of Economics; *healy.52@osu.edu*

[§]California Institute of Technology, Division of the Humanities & Social Sciences; *kirby@caltech.edu*

I. INTRODUCTION

A growing literature documents that choices are often deliberately stochastic, a phenomena referred to as *random choice*. The desire to randomize is prevalent in many domains and has been documented using different methodologies (Agranov and Ortoleva, 2022). However, what we do not know is whether choice stochasticity is correlated across domains for a given individual. In other words, is randomization itself a stable trait? Are those who randomize in one domain more likely to randomize in another? For example, can mixing over lottery choices predict mixing in games?

To shed light on these questions, we measure individual propensity to randomize in different domains and explore the cross-domain correlations. We study four domains: (1) *Probability Matching* (PM) problems, which involve the choice between first-order stochastically dominating and dominated options, (2) *Risky-Safe* (RS) problems, which involve choices between a risky lottery and a sure amount, (3) *Strategic Certainty Games* (SC), which are isomorphic to the PM problems but framed as subjects playing a 2×2 game against a known distribution of opponents' actions, and (4) *Strategic Uncertainty Games* (SU) in which subjects play a 2×2 game against other participants in the current session. For each decision problem in each domain, we ask the subject to make their choice twenty times, with one of the twenty being randomly selected for payment. This is our baseline treatment, which we call IND, indicating that the twenty repetitions correspond to twenty independent realizations of uncertainty, i.e., twenty different draws of a ball from an urn or twenty different opponents in games.

In all four domains we see that a substantial fraction of subjects randomize by varying their choices over the twenty repetitions. Nearly 70% of subjects mix in at least one decision problem in each domain, including the PM problems in which one option stochastically dominates the other. More importantly, we find strong correlation across domains: 52% of subjects mix in all four domains and 17% never mix in any domain. This indicates that there are “mixing types,” who have a preference for randomizing in all of the domains, and “non-mixers” who always pick the same option in all twenty repetitions.

In three auxiliary treatments, we explore the robustness of these mixing types. First, we study whether subjects randomize because they incorrectly believe in negative serial correlation in the outcomes of the twenty lottery choices. This belief could manifest in randomization since negative serial correlation implies that the optimal

choice is varying across the twenty repetitions. Our CORR treatment eliminates this possibility by having all twenty choices pay based on one single draw from the Bingo cage (or, the choice of one single opponent). Thus, if Option A pays more in one choice, it will pay more in all twenty choices. Surprisingly, we find that mixing behavior in CORR is indistinguishable from the IND treatment, suggesting that mixing behavior cannot be rationalized by subjects having different beliefs about the twenty replicated choices. We discuss other possible theories and heuristics in Section VI.

Second, we examine the robustness of mixing types by having subjects condition on each choice as though it were the choice that will be paid. Specifically, in the SEQ treatment, we ask subjects to make their twenty choices sequentially, learning after each replicate whether or not it is paid before moving to the next. Once subjects learn that a particular replicate was paid, they stop and move on to the next decision problem.¹ This encourages subjects to think of the current choice in isolation, rather than as part of a portfolio of twenty. We find that the percentage who mix in the PM questions (where one option dominates the other) drops by over 20 percentage points to around 40%, but does not change in the RS questions. This shows that there is a mixing type that is responsive to interventions. We then ask these same subjects to participate in the IND treatment, and find that mixing in the PM question remains at 40%. Thus, the responsive mixers appear to learn from the SEQ treatment that mixing is not what they prefer, and continue not to mix in our original treatment.

Finally, we run an online replication of our study to test whether experimenter demand effects may be driving our results. Following de Quidt et al. (2018), we introduce a new treatment in which the instructions strongly suggest that the researchers want subjects *not* to mix. We find that the percentage of subjects who mix is unaffected by this intervention, suggesting that mixing is robust to experimenter demand effects. The responses in the non-incentivized open-ended questionnaire conducted at the end of the experiment confirm that most subjects randomize consciously and provide coherent reasons for mixing.

One way to assess the importance of understanding randomization behavior is through financial outcomes. We find that randomization leads to lower expected earnings when one of the options is dominated. But when comparing a risky alternative to a sure payment, randomization doesn't significantly affect earnings. Thus, the

¹Thus, when making their twelfth choice (for example) they know that the first eleven were not paid, and that they will not face choices thirteen through twenty if the twelfth choice is paid.

financial costs or benefits of randomization can vary across settings. Interestingly, we find that subjects who don't mix in the SEQ treatment are less likely to mix when returning back to the original IND design. Apparently they learn not to mix from the SEQ treatment. But this only occurs for PM questions (where one option is dominated); in the RS questions we see no significant drop in mixing or evidence of learning. Thus, some subjects may learn that mixing over a dominated option is a mistake, and this can improve their outcomes. But many mixers do not exhibit any learning. For these subjects mixing appears to be a very stable preference.

Taken together, our results suggest that randomization is an individual trait, which is stable and predictable across domains. This finding has a few important implications outside of a stylized laboratory environment. First, there is field evidence that individuals randomize in important high-stakes decisions. For example, Dwenger et al. (2018) show that individuals randomize their reported college rankings in strategy-proof mechanisms, Zhang and Zhong (2020) find that individuals randomize to resolve choice overload in charitable donations, and Levitt (2021) shows that individuals are willing to flip a coin to make major life decisions. Our results suggest that these behaviors could be manifestation of fundamental individual-level preferences for randomization, rather than exploration or exceptional behavior of a few individuals.

II. RELATION TO THE LITERATURE AND CONTRIBUTION

Our results most obviously inform theories of individual decision-making. In Section VI we discuss several natural behavioral models that rely on subjects' ability to reduce compound lotteries and show that no such theory can rationalize our data. We then consider theories that do not assume reduction, but instead define preferences over two-stage lotteries, and identify a class of perturbed utility models that can predict mixing. While they can organize the data, they are a bit unsatisfactory in that they do not provide an underlying rationale for randomization; the model predicts that people mix because they have a preference for mixing. Finally, we explore various heuristics and biases and find that these are also unable to explain our data. We therefore conclude that the mixing we observe represents a real challenge to existing theories, and that developing theories or heuristics that can accommodate this behavior would be a fruitful avenue for future research.

The correlation across domains also suggests that we should also consider prefer-

ences for randomization in other contexts. For example, we can view our results as providing a foundation for non-equilibrium mixing in games.² Many game-theoretic concepts are built around the idea that players tremble, or that they best respond with noise. Quantal response equilibrium (McKelvey and Palfrey, 1995) is a solution concept that explicitly incorporates noisy best response. This noise is often modeled as arising from payoff shocks or misspecifications, but one could alternatively interpret it as reflecting deliberate randomization by players who simply prefer to mix (Allen and Rehbeck, 2021). Indeed, we find that our mixing subjects put more weight on a lottery or strategy when its expected value increases, consistent with the common assumption that better responses are played more frequently (Goeree et al., 2005).

Our results add a new branch to the literature on stability of behavioral types across domains.³ Most studies on risk preferences find that they are not stable (Binswanger, 1980; Isaac and James, 2000; Kruse and Thompson, 2003; Eckel and Wilson, 2004; Berg et al., 2005; Anderson and Mellor, 2009; Vlaev et al., 2009; Dulleck et al., 2015) except when the domains are very similar (Choi et al., 2007; Reynaud and Couture, 2012; Slovic, 1972). Time preferences appear to be stable across different delay lengths (McLeish and Oxoby, 2007; Halevy, 2015) and do not seem to depend on which good is being consumed (Reuben et al., 2010; Ubfal, 2016). Most studies find fairly consistent patterns of social preferences across settings (Fisman et al., 2007; Ackert et al., 2011; De Oliveira et al., 2012b), though there are exceptions (Blanco et al., 2011). Cross-game correlations of strategic sophistication seem more mixed, with high rates of correlations in some families of games, but not others (Georganas et al., 2015). Our results suggest that mixing is one type that is stable across domains, and the fact that individuals mix across domains may also help explain the fact that other traits—such as risk aversion and strategic sophistication—appear to be unstable across decision problems.

²Mixing in normal-form games can be rational when the player is exactly indifferent between strategies. In extensive-form games mixing may be used to discipline the behavior of another player, for example when a tax authority chooses an optimal auditing policy (Becker, 1968). We study mixing only in normal-form games and our design rules out indifference as an explanation.

³A related literature studies stability over time for risk preferences (Horowitz, 1992; Hey, 2001; Harrison et al., 2005; Dave et al., 2010; Reynaud and Couture, 2012; Crosetto and Filippin, 2016), time preferences (Kirby, 2009; Krupka and Stephens, 2013; Meier and Sprenger, 2014), and social preferences (Brosig et al., 2007; De Oliveira et al., 2012b,a; Lotz et al., 2013; Lönnqvist et al., 2015; Bruhin et al., 2019). Chuang and Schechter (2015) provide a survey, concluding that experimental measures of social and time preferences show weak intertemporal correlation, while risk preferences exhibit zero correlation.

Finally, our paper relates to a large experimental literature on randomization, most of which focuses exclusively on a single domain. Papers have studied probability matching over stochastically dominant and dominated lotteries (Humphreys, 1939; Grant et al., 1951; Siegel and Goldstein, 1959; Loomes, 1998; Rubinstein, 2002), randomization in decisions that do not feature dominant options (Sopher and Narramore, 2000; Dwenger et al., 2018; Agranov and Ortoleva, 2017; Feldman and Rehbeck, 2019; Agranov and Ortoleva, 2022), and randomization in games (Shachat, 2002; Sandroni et al., 2013; Romero and Rosokha, 2019). First, we complement this literature in showing similar rates of randomization at a subject-level. These papers vary widely in terms of the decisions used, but our results fall within the established frequencies of randomization. For example, Dwenger et al. (2018) find that about a third of subjects randomize over consumption goods or college rank-order reports. In more similar domains, Agranov and Ortoleva (2017) find 71% of subjects change their response to lottery questions repeated in a row, and Feldman and Rehbeck (2019) find that 94% of subjects randomize in the domain of risk. We find that between 64% and 76% subjects mix in at least one question in each of our domains, commensurate with the previous literature.

In addition to documenting similar patterns, our paper is the first experiment to show within-person correlations in mixing behavior across these different domains. We find strong evidence of “mixing types” who randomize in all environments, and provide evidence that some of the observed dominated mixing is responsive to intervention. To our knowledge, our experiment is also the first to study randomization behavior in games compared to their equivalent decision problems, and to show evidence of mixing in games that cannot be rationalized.

III. DESIGN OF THE INDEPENDENT EXPERIMENT

Our main experiment is called the Independent (IND) treatment. It consisted of four sessions with 21 subjects in each for a total of 84 subjects and it was conducted at the Ohio State Experimental Economics Laboratory.⁴ We used physical randomization devices—draws from a bingo cage and rolls of dice—to resolve all uncertainty.

⁴Subjects were recruited through ORSEE (Greiner, 2015). No subject participated in more than one experimental session. The software was custom-built using PHP and MySQL. Subjects interacted with the software via a web browser on private computer terminals. Sessions lasted roughly 90 minutes, and subjects earned on average \$22.41 (which includes a \$5 show-up fee).

Each experimental session consisted of four decision blocks, with each block comprising a different type of decision task. The order of blocks was randomized across sessions, with the only restriction being that Block IV (risk elicitation) always appeared last. Instructions for each block were distributed and read out loud to subjects before the start of the block. In addition, the experimenter used slides as a visual aid to clarify the procedures and the tasks. We paid subjects for one randomly-selected choice made in the experiment. First, we describe the decisions that subjects faced in each block, and then we describe the payment procedure in detail at the end of this section. The complete instructions appear in the Appendix. Many of the choices in the experiment involve lotteries with two outcomes. We write $(\$a, p; \$b)$ to denote a lottery that pays $\$a$ with probability p and $\$b$ with probability $1 - p$; $(\$a, 1)$ denotes the degenerate lottery that pays $\$a$ with certainty.

Block I: Individual Decisions. This block consisted of twelve questions: six questions that involve first-order stochastically dominated options, which we refer to as *probability matching* (PM) questions, and six choices between a risky lottery and a sure amount, which we refer to as *risky-safe* (RS) questions. These twelve questions were presented to subjects in random order, each question on a different screen. In each of the twelve questions, a subject chose between the same two lotteries twenty times, all on the same screen. In PM questions, each decision involved choosing twenty times between a dominant bet of $(\$25, p; \$5)$ with $p > 1/2$, and a dominated bet, $(\$25, 1 - p; \$5)$. In RS questions, each decision involved choosing twenty times between a risky bet, $(\$25, p; \$5)$ (again with $p > 1/2$), and a safe bet, $(\$15, 1)$. The six PM questions differed only in the probability associated with the dominant bet, i.e., $p \in \{0.55, 0.60, 0.65, 0.70, 0.75, 0.80\}$. Similarly, the six RS questions differed only in the probability associated with the risky bet, where $p \in \{0.55, 0.60, 0.65, 0.70, 0.75, 0.80\}$. We refer to these questions by their acronym and associated probability, e.g. PM55 refers to the twenty choices between $(\$25, 0.55; \$5)$ and $(\$25, 0.45; \$5)$, and RS55 refers to the twenty choices between $(\$25, 0.55; \$5)$ and $(\$15, 1)$.

Both RS and PM questions were presented in terms of betting on a ball drawn from a Bingo cage. We had a Bingo cage filled with twenty balls, numbered 1–20. Each bet specified payoffs that a subject would receive depending on which ball would be drawn from the cage. For example, in the RS questions, choosing, say, the risky bet $(\$25, 0.75; \$5)$ indicates winning $\$25$ if the ball drawn is numbered 1–15 and

winning \$5 if the ball drawn is numbered 16–20, while choosing the safe bet (\$15, 1) indicates winning \$15 regardless of which ball is drawn. Similarly, in PM questions, choosing the dominant bet (\$25, 0.75; \$5) indicates winning \$25 if the ball drawn is numbered 1–15 and winning \$5 if the ball drawn is numbered 16–20, while choosing the dominated bet (\$25, 0.25; \$5) indicates winning \$5 if the ball drawn is numbered 1–15 and winning \$25 if the ball drawn is numbered 16–20.⁵ Which of the twenty bets would be chosen for payment was determined by the roll of a twenty-sided die.

	Left	Right		Left	Right
Up	\$25, \$5	\$5, \$25	Up	\$25, \$15	\$5, \$5
Down	\$5, \$25	\$25, \$5	Down	\$15, \$25	\$15, \$15
	<i>Matching Pennies</i>			<i>Dominance Solvable</i>	

Table I: Matching Pennies and Dominance Solvable Games

Block II: Games with Strategic Certainty. In this block, subjects played a Matching Pennies game twice, both times against a known distribution of past players' actions. We reproduce the payoff matrix for the game on the left side of Figure I above. Subjects were presented with the game form as the row player, and told that they would be matched randomly with one of twenty column players who had played the game in a previous session.⁶ We tell subjects the truthful distribution of past players who played Left and Right. In one iteration of the game, we tell subjects that 11 of 20 (55%) previous players played LEFT, and in the other we tell them that 16 out of 20 (80%) previous players played LEFT. We refer to these games as SC55 and SC80, respectively. While framed very differently, information about the distribution of actions of past players creates isomorphism between these games with strategic certainty and two of the PM questions from Block I (PM55 is isomorphic to SC55 and PM80 is isomorphic to SC80), where playing UP corresponds to choosing the dominant bet (\$25, p ; \$5). Subjects chose between Up and Down twenty times for both distributions of past players.

These games with strategic certainty allow us to fix subjects' beliefs in comparing decision problems to games. If we find differences in mixing behavior between PM and

⁵Given this structure of payoffs, the dominant bet is first-order stochastically dominant but is never state-wise dominant.

⁶The data for the past players was collected at University of California in Irvine in December 2017.

SC questions, this suggests that the mere *framing* as a game affects randomization tendencies. This might indicate that individuals do not treat uncertainty from nature in the same way they treat uncertainty from other individuals, indicating that mixing types are sensitive to the nature of uncertainty.

Block III: Games with Strategic Uncertainty. Subjects played two different 2×2 matrix games, Matching Pennies and a Dominance Solvable game, against current opponents in the room. We refer to these games as SUMP (*Strategic Uncertainty Matching Pennies*) and SUDS (*Strategic Uncertainty Dominance Solvable*), respectively. As described above, the Matching Pennies game is equivalent to a PM question for a given belief of their opponent choosing LEFT. The Dominance Solvable game is equivalent to a RS question, where choosing UP corresponds to choosing the risky option and choosing DOWN corresponds to choosing the safe option. Subjects played through each game in five different stages. Within each game, the stages were presented in the order described below, and the order of the two games was randomized across sessions.

In Stage 1, subjects played the game for a single repetition as Column player, choosing either LEFT or RIGHT. In Stage 2, subjects played twenty repetitions of the game as Row player, all on the same screen, just as described in Block II above. Each of their twenty row choices could be matched with a random Column player's decision from Stage 1. In Stage 3, we elicit subjects' belief that a random Column player chose LEFT. We use these beliefs to compare games with strategic uncertainty to their analogous Block 1 individual-choice questions.⁷

It could be that individuals mix in these games because their beliefs are such that they are exactly indifferent between UP and DOWN. This is especially plausible in the Matching Pennies game. We address this possibility in Stages 4 and 5 by giving subjects a noisy signal of one opponent's play—which should cause a change in their beliefs away from indifference—and ask them to play the game again. Specifically, in Stage 4 subjects see a signal of one opponent's action, LEFT or RIGHT, that is correct 55% of the time but incorrect 45% of the time. Subjects then play another 20 repetitions of the game as Row player, just as in Stage 2. If they were mixing in Stage

⁷We required subjects to report a point belief in order to match games to individual decision problems. Recent work suggests that individuals might have stochastic beliefs in games (Friedman and Ward, 2022, e.g.) which we cannot capture here.

2 due to indifference, we should not see any mixing in Stage 4 (and vice-versa). In Stage 5, we elicit post-signal beliefs, just as in Stage 3. We will say a subject mixes if they mix in *both* Stage 2 and Stage 4.

Block IV: Risk Elicitation. Subjects complete two standard risky investment tasks (Gneezy and Potters (1997), Charness et al. (2013)) presented in a random order to measure their risk preferences. In each task, subjects are endowed with \$10, any portion of which they could invest in a risky project. If the project is successful, which occurs with probability p , the amount invested is multiplied by R and paid to the subject. If the project is unsuccessful, the amount invested is lost. In either case, subjects keep the portion of the endowment they chose not to invest. The parameters used in the two tasks are $(p = 0.5, R = 2.5)$ and $(p = 0.4, R = 3)$. This risk elicitation method is popular due to its simplicity and time requirements. Moreover, administering it twice has been shown to reduce measurement error associated with elicitation of risk attitudes (Gillen et al. (2019)).

To summarize, subjects go through 26 decision problems during the session. 14 are individual decisions (Blocks I and IV) and the remaining 12 are game decisions (Blocks II and III). Some of the questions have one repetition (both risk attitude questions in Block IV and Stages 1, 3, and 5 of games with strategic uncertainty in Block III), while all the remaining questions have twenty repetitions of the same choice presented on the same screen.

Subjects' payments. At the end of the experiment, one of the questions was randomly selected for payment.⁸ The same question was selected for all subjects in a session, but the selected question could differ between sessions. If the selected question had only one repetition, then subjects were paid based on this single choice. If the selected question had twenty repetitions, then we used the physical Bingo cage and dice to determine subjects' payments. Specifically, the experimenter had a transparent Bingo cage filled with twenty balls numbered from 1 to 20. First, the experimenter drew twenty balls with replacement and wrote these draws on the board, so that all subjects observed these draws. Each of these draws corresponded to one of the twenty repetitions of a choice in the selected question. After the twenty draws were recorded,

⁸Each question was equally likely to be selected for payment. See the online appendix for a discussion of incentive compatibility of this mechanism in our setting.

the experimenter rolled a 20-sided die to determine which of the twenty repetitions would be selected for payment. For example, if the die came up 17 and the 17th ball drawn was ball 5, then we look at the subject’s choice on the 17th repetition and pay the bet chosen based on ball 5 being drawn. Using two different randomization devices emphasizes to subjects that the choice of which repetition is paid is unrelated to whether the chosen bet pays off.⁹ We paid all subjects for the same repetition of a question. Subjects observed choices they submitted in the selected question during this “theatrical” performance of the experimenter.

IV. INDEPENDENT EXPERIMENT: RESULTS

Our main object of interest is the tendency of subjects to randomize their answers across repetitions within a decision problem. This requires a definition of randomization at the individual level. In the analysis that follows, we identify a subject as a *mixer* if they chose strictly less than 90% of same bets in a decision problem, i.e., fewer than 18 same bets out of 20 total repetitions of the same choice. We identify a subject as mixing in a given domain if they were a mixer in at least one of the questions in that domain. In Appendix B, we show that while levels of mixing are obviously responsive to this cutoff, the qualitative results remain the same.

One could imagine a definition of mixing based on the intensive margin. For example, we could count the number of times out of twenty that the subject chooses Option A. But theories that rationalize mixing (see Section VI) predict that the number of times Option A is chosen varies from one problem to the next, based on the underlying expected utility of the two options. Thus, such a definition would not be comparable across problems or across domains without specifying a particular functional form. Because of this, we take the null hypothesis to be that a subject does not randomize, and we define a mixer as someone who meaningfully deviates from that null. For the interested reader, histograms of the number of times each option was chosen are provided in Appendix II.D.

⁹If a Strategic Uncertainty game is paid (Stages 2 and 4 of Block III) then one subject is randomly chosen as the row player and another subject is randomly chosen as the column player, and they are paid based on their strategy choices. Specifically, one of the row player’s 20 repetition is randomly selected and compared to the column player’s Stage-1 choice. All other subjects are paid a flat payment of \$15. If a belief elicitation task is paid (Stages 3 and 5) then all subjects are paid according to the BDM-for-probabilities method (Grether, 1981), which is incentive compatible as long as subjects’ preferences respect dominance; see Online Appendix A for details.

In all regression analyses we cluster standard errors at the individual level to avoid interdependencies of observations that come from the same subject completing several tasks in the experiment. Bar graphs are shown with 95% confidence intervals.

IV.A. Mixing Types

We classify subjects into three mutually exclusive types based on their behavior in all domains: Subjects who do not mix in any of the domains are called *Never Mix* (17% of subjects), subjects who mix in all domains are called *Always Mix* (52% of subjects), and subjects who mix sometimes but never violate first-order stochastic dominance, i.e., mix in PM questions or SC games, are called *Non-Dominated Mix* (2% of subjects). Thus, the modal subject randomizes in every environment we considered.

One unexpected pattern is that 8% of subjects mix in games (SC and SU) but not in decision problems (PM or RS). All other patterns are rare, each explaining at most 3.6% of subjects. Overall, it is most common for subjects to *always* mix or *never* mix. A minority of subjects mix in some environments but not in others. Thus, randomization behavior appears to be a “type” that varies among individuals but is generally consistent across environments.

Result 1. *Randomization is an individual trait. The two most prominent types in the population are subjects who Never Mix (17%) and those who Always Mix (52%).*

To further investigate individual types, Table II presents pairwise correlations for each pair of domains. We find that mixing is positively and significantly correlated across all domains. We see strong correlations across different decision environments, suggesting that individual types are robust and not solely driven by decision framing.

	PM	RS	SC	SUMP
RS	0.71***			
SC	0.52***	0.48***		
SUMP	0.58***	0.54***	0.71***	
SUDS	0.42***	0.35***	0.37***	0.46***

Table II: Pairwise Correlations in Individual Mixing in the IND Experiment

Notes: We report pairwise correlations between indicator variables indicating whether a subject mixed in each of our decision environments. *** indicates significance at 1% level.

However, we do find some differences that would be interesting to explore in future

work. Individuals appear to mix more in games than in their equivalent decision problems: 25% of individuals who mix in SC55 do not mix in PM55 compared to only 8% of mixers in PM55 who do not mix in SC55. This is similar with 51% of mixers in SC80 mixing in PM80, whereas 83% of mixers in PM80 mix in SC80. Further, we compare behavior in the SUMP decisions to those in PM questions for a subsample of subjects for whom we are able to match these decisions (68% of subjects).¹⁰ For these subjects, we find that they are more likely to mix in SUMP than in the corresponding PM decision problem (68% mixers in SUMP vs. 49% in PM, signed-ranks test $p < 0.001$), and they choose significantly more dominated bets in SUMP than in PM (84% dominant bets in PM vs. 70% dominant actions in SUMP, signed-ranks test $p < 0.001$).¹¹

Result 2. *Mixing is highly correlated across domains at an individual level. In addition, individuals are more likely to mix in games than in their equivalent individual choice problems, for both games with strategic certainty and strategic uncertainty.*

IV.B. Variations Within A Domain

While we find that mixing is a stable type *across* domains, the prevalence of mixing also responds sensibly to changes *within* a domain. Figure I depicts the frequency

¹⁰To compare games with strategic uncertainty to corresponding decision problems, we match individuals' beliefs in SUMP (SUDS) to the corresponding objective probabilities in PM (RS) decision problems. We can match the SUMP game for an individual with belief $p(LEFT) = 0.55$ or $p(RIGHT) = 0.55$, for example, to the PM55 question. For belief $p(LEFT)=0.55$, the dominant action is UP, whereas for $P(RIGHT)=0.55$, the dominant action is DOWN. Both situations are isomorphic to the PM55 decision problem. In the SUDS game, an individual with belief e.g. $p(LEFT) = 0.80$ can be matched with the RS80 decision problem. We focus on subjects' beliefs after the signal, as this is where individuals are less likely to have 50–50 beliefs in the Matching Pennies game. 26% of individuals have 50–50 beliefs before the signal, and only 11% have 50–50 beliefs after the signal. Recall that we consider a subject to be a mixer only if they mix both before and after the signal.

¹¹We can only match SUDS to RS questions for 29% of subjects, so our comparison sample sizes are much smaller for this game. Given that the column player has a dominant strategy, it is not surprising that we match fewer subjects in SUDS. About 60% of subjects have a belief higher than 80%. These subjects are directionally more likely to mix in SUDS than in the equivalent RS decision problem (54% mixers in SUDS vs. 33% in RS, signed-ranks test $p = 0.18$), and choose more risky actions in the SUDS than in RS (51% risky actions in SUDS vs. 36% in RS, signed-ranks test $p = 0.18$). Reported results are for exact matching of beliefs. If we round beliefs to the nearest 5, we can match 77% of subjects in SUMP and 37% of subjects in SUDS. Qualitatively, the results are the same. We find more mixers in SUMP than in the corresponding PM questions (72% in SUMP vs 52% in PM, $p < 0.001$), and we find they choose more dominant bets in PM than in SUMP (82% in PM vs. 69% in SUDS, $p < 0.001$). We find more mixers in SUDS than in the corresponding RS questions (58% in SUDS vs 42% in RS, $p = 0.18$), and we find they choose more risky bets in SUDS than in RS (47% in SUDS vs. 35% in RS, $p = 0.16$).

of mixing in each question as well as overall mixing in each domain (the last four bars on the far right part of the figure). Mixing behavior is very common in all four domains: between 64% and 76% subjects mix in at least one question in every single domain. The frequencies of mixing are quite similar across domains except for the slightly higher frequency detected in the games with strategic certainty (SC).¹²

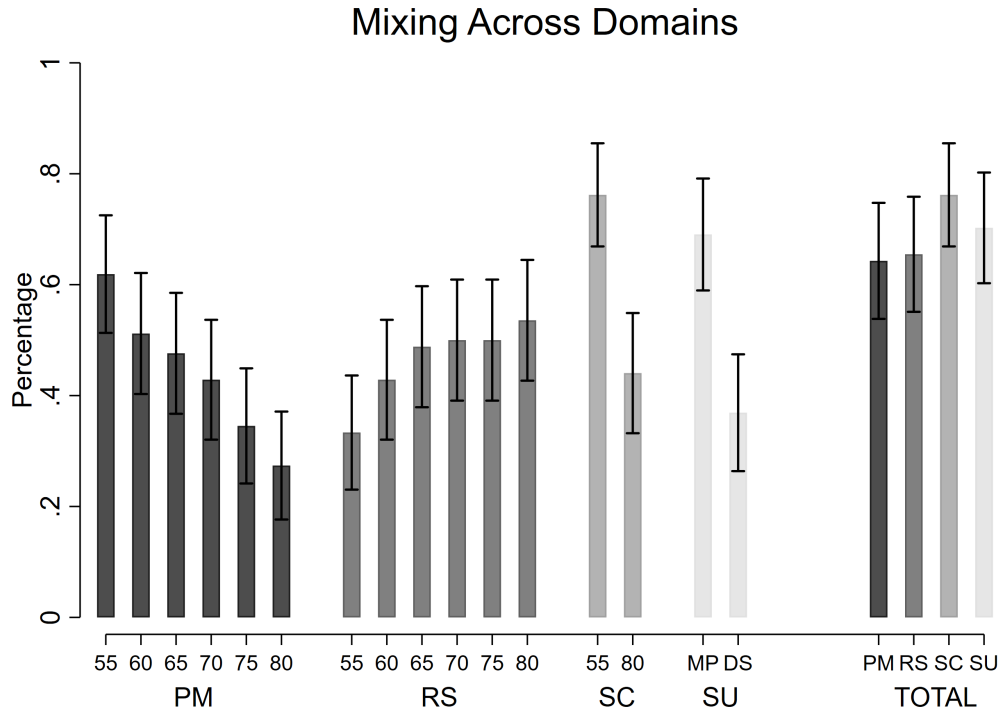


Figure I: Percentage of Subjects Who Randomize in Each Question or Domain

Notes: The error bars depict 95% confidence intervals.

Within the PM domain, the chance of the dominant bet paying off ranges from 55% to 80% across questions. We observe that subjects react to this change in probability in a monotone manner: as the dominant bet becomes “more dominant,” individuals become less likely to mix (a Probit regression coefficient, -0.064, is negative with $p < 0.001$), and those who do mix choose the dominated bet less often (linear regression, $p < 0.001$). Those who mix choose the dominant bet in only 11.6 of 20

¹²We ran a Wilcoxon signed-rank test to compare frequency of mixing between each pair of domains: $p = 0.763$ PM vs. RS, $p = 0.018$ PM vs. SC, $p = 0.197$ PM vs. SU, $p = 0.039$ RS vs. SC, $p = 0.317$ RS vs. SU, and $p = 0.096$ for SC vs. SU

repetitions in the PM55 question, but choose it in 15.5 of 20 repetitions in PM80.¹³

Similar analysis for RS questions reveals that subjects react to the probability of the risky bet paying off, which ranges in RS questions from 55% to 80%. In the most risky question (55%) around 2/3 of subjects are non-mixers, and the vast majority of them are choosing only the safe bet. As the risky bet becomes less risky, subjects become more likely to mix by adding in the risky bet. The marginal effect on the percentage bet variable in the Probit regression is estimated at 0.035 ($p < 0.001$). Moreover, as the risky bet becomes less risky, subjects who do mix choose it more often with an average of 6.1 out of 20 repetitions in RS55 question and 9.8 out of 20 repetitions in RS80 question. Regression analysis confirms this: the estimated coefficient on the percentage risky variable is 0.749 ($p < 0.001$).

The fact that mixing in PM and RS domains varies with the probabilities informs the way we think about the SC and SU games. In the SC games, most of the mixing comes from SC55, in which the distribution of choices of past players is close to uniform (76% of subjects mix in this case), while only 44% mix in SC80 in which the distribution of past players' choices has much smaller variance. Similarly, most of the mixing in SU games happens in the Matching Pennies game (69% of subjects mix in this game), while less than 40% do so in the Dominance Solvable game.¹⁴

Result 3. *Randomization is highly prevalent in all domains but responds sensibly to the available lotteries.*

V. ADDITIONAL EXPERIMENTS

We conduct three follow-up experiments to test the robustness of mixing types along three dimensions. The first, our Correlated Experiment, tests whether mixing types are sensitive to the nature of uncertainty. The second, our Sequential Experiment, tests whether mixing types are robust to manipulations of contingent reasoning. The third, our Online experiment, addresses the concern that the mixing is driven by an experimenter demand effect by utilizing de Quidt et al. (2018) method. In this

¹³The mixing we observe in our PM questions does not necessarily contradict the findings of Agranov and Ortoleva (2017), who see no mixing when one lottery dominates the other state-by-state. Our PM questions feature stochastic dominance, but not state-wise dominance; therefore, we conjecture that subjects mix with stochastic dominance, but not with state-wise dominance.

¹⁴Statistical analysis confirms that the fraction of mixing behavior in SC55 game is significantly higher than that in SC80 game (signed-ranks test $p < 0.001$). Similarly, significantly more subjects mix in SUMP game compared with SUDS game ($p < 0.001$).

section, we first discuss each of the three additional experiments. Second, we compare risk attitudes of different mixing types (elicited in the CORR and IND treatments) to highlight how mixers and non-mixers can differ ‘out-of-sample’.

V.A. *The Correlated Experiment*

Prelude. One possible explanation for the “*Always Mix*” types observed in the IND experiment is a “gambler’s fallacy” belief where subjects incorrectly expect negative serial correlation across the twenty independent draws from the Bingo cage. To study whether this is indeed what drives randomization behavior in our IND experiment, we conducted the second experiment called the Correlated Experiment (CORR).

Design. The CORR experiment has the exact same structure as the IND treatment including the composition of the Bingo cage and the descriptions of all tasks. However, if the selected question contained twenty repetitions, then instead of drawing twenty balls, *only one* ball was drawn from the Bingo cage. After that, the experimenter rolled the 20-sided die to determine which repetition would be paid. That is, a subject in the CORR experiment knew that regardless of which of their twenty bets were chosen for payment, they would all pay out against the exact same ball drawn from the Bingo cage.

Therefore, the IND and CORR treatments were identical except for the realizations of uncertainty. In the IND experiment, each of the subject’s twenty decisions corresponded to a different independent realization of uncertainty. This means that the ex-post optimal bet could differ across the twenty decisions. In the CORR experiment, however, each decision corresponded to the same single realization of uncertainty. This means that the ex-post optimal bet is the same for all twenty decisions by construction. Thus, the CORR treatment minimizes potential misconception that subjects might have about serial auto correlation between realizations of uncertainty. For example, subjects in the IND experiment could believe a high numbered ball is “due” after a low numbered ball, which might cause them to alternate which bets they chose. In the CORR experiment, however, only *one* ball is chosen, so subjects should not hold such a belief. If this is indeed the underlying reason for mixing behavior in any of the domains, then we expect to see less mixing in CORR than in the IND Experiment.

84 new subjects participated in the CORR experiment, which was also conducted at the Ohio State Experimental Economics Laboratory.

	IND Experiment	CORR Experiment
Never Mix	17%	13%
Always Mix	52%	45%
Non-Dominated Mix	2%	11%
Others	29%	31%
# of subjects	84	84

Table III: Individual Types in the IND and CORR Experiments

Notes: Subjects who do not mix in any of the domains are called *Never Mix*, subjects who mix in all domains are called *Always Mix*, and subjects who mix sometimes but never violate first-order stochastic dominance, i.e., mix in PM questions or SC games, are called *Non-Dominated Mix*.

Results. The classification of subjects into individual types yields similar results to those obtained in the IND Experiment (summarized in Table III). The two most prevalent types remain subjects who mix in all domains (45% in the CORR experiment) and those who never mix (13% in the CORR experiment). Within the category of “Others,” we can identify two predominant patterns: 10% mix in all the domains *except* for games with strategic uncertainty and 11% mix in all domains except for RS questions. The Fisher exact test and the chi-squared test comparing the distribution of types in the IND and CORR treatments show that these distributions are not statistically different ($p = 0.15$). Therefore, it appears that the nature of uncertainty does not affect subjects’ mixing type.

Furthermore, we find that this correlated uncertainty does not affect mixing behavior in any of our domains. Figure II compares the frequency of mixing in the IND and CORR experiments within each domain. We find no significant differences ($p > 0.51$ for all pairwise domain comparisons, $p > 0.27$ for question-specific comparisons).

Therefore, incorrect belief in serial correlation is *not* the driving force of mixing behavior, nor a determinant of subjects’ mixing type. Regardless of whether the realization of uncertainty happens twenty times independently (as in the IND Experiment), or once for all twenty choice (as in CORR Experiment), individuals mix to the same extent. This also allows us to rule out a number of potential explanations for randomization behavior, as we discuss in Section VI.

Result 4. *Mixing types are robust to correlated realizations of uncertainty, suggesting that mixing behavior is not driven by a mistaken belief in the negative serial correlation in draws.*

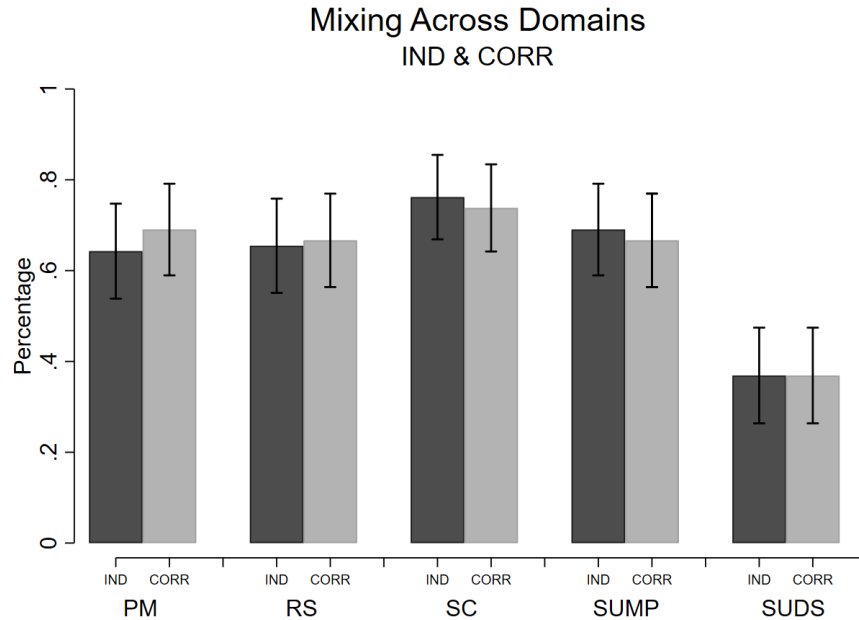


Figure II: Mixing Behavior in the IND and CORR Experiments

Notes: The error bars depict 95% confidence intervals.

V.B. The Sequential Experiment

Prelude. The Sequential Experiment investigates whether mixing in the PM and RS questions is driven, in part, by a failure of contingent reasoning (Esponda and Vespa, 2019, e.g.). While individuals might be able to recognize dominance in a single decision, they might fail to think contingently when answering the same question multiple times. We conjecture that individuals will be more likely to choose all dominant bets when they are encouraged to treat each of the twenty choice repetitions “in isolation,” which will affect our individuals classified as “always mix” types. In the RS questions, however, treating each decision in isolation need not result in less mixing as the subject might not clearly prefer one option over another. Thus, our Sequential Experiment tests the robustness of mixing types to an environment that encourages contingent reasoning and isolation of each choice.

Design. The Sequential experiment consists of three within-subject treatments: Sequential (SEQ), Simultaneous (SIM), and One (ONE). Subjects first participate in the SEQ block, then in the SIM block, and then in the ONE block. In each block,

subjects face the same twelve questions (six PM and six RS) as in the IND experiment, presented in random order. We did not include the SC or SU games in this experiment.

The goal of the SEQ treatment is to encourage subjects to treat each of the twenty repetitions as a single unique choice. At the beginning of each decision problem, the computer randomly selected which of the twenty repetitions would be the one chosen for payment if that decision problem were selected. Subjects made their choice in each of the repetitions sequentially. After each decision was recorded, a subject learned whether that repetition was the one chosen for payment. If it was not, the subject moved onto the next repetition. If it was, the decision problem terminated and the subject moved on to the next decision problem without answering the remaining repetitions.¹⁵ This treatment encouraged subjects to treat each of the twenty repetitions as if it were the only choice to be made. While making their current decision, subjects knew that any previous repetitions certainly would not be paid, so the current choice was made in isolation.

The second environment was the Simultaneous (SIM) treatment, which was exactly the same as our IND Experiment; that is, subjects made twenty choices simultaneously in each decision problem. The only difference between the SIM treatment and the IND Experiment was that subjects in the SIM treatment just previously had participated in the SEQ treatment. If they “learned” contingent reasoning by participating in the SEQ treatment, we could see less mixing in SIM than in the IND Experiment.

The last environment was the ONE treatment, where individuals made each binary choice only *one* time. Subjects saw each of the PM and RS questions in random order and chose their preferred bet one time for each question. They only saw one choice repetition on their screens for each question. The goal of this treatment was to establish individuals’ “isolated” preference on a given decision problem. We expected that most individuals would choose the dominant bet in all PM questions, but would choose either the risky or safe option in RS questions according to their risk preferences.

¹⁵For example, imagine the computer selected the sixth repetition as the one chosen for payment. The subject makes their choice on the first repetition, then learns this was not the one chosen for payment. Then they make their choice on the second repetition, and again learns it was not chosen for payment, and so on. After making their choice on the sixth repetition, they learn that this was the one chosen for payment. Then, the subject moves onto the next decision screen and never answers repetitions seven through twenty. As the subject makes each of these sequential decisions, they see all twenty repetitions on their screen as before, but the subsequent repetitions are greyed out until they actually makes the decision.

We also conducted Sequential Experiment at the Ohio State University Experimental Economics Laboratory with 93 new participants.

Results. The behavior in the ONE treatment confirms our expectations: When subjects are asked to make only one choice between the dominant and the dominated bets in the PM questions, they almost always pick the dominant one. At the same time, when subjects make only one choice between the risky and the safe bet in the RS questions, their choice responds sensibly to the probability of getting a high prize in the risky bet. Table IV makes this point by showing that more than 90% of subjects in the PM questions in ONE treatment select the dominant bet irrespective of the likelihood of getting the high prize in this dominant bet, while the fraction of subjects who choose the risky lottery in the RS questions increases monotonically as the risky lottery becomes more and more attractive.

	<i>Probability of Bet</i>					
	55%	60%	65%	70%	75%	80%
ONE						
% Dominant (PM)	91%	94%	98%	98%	99%	98%
% Risky (RS)	11%	19%	25%	45%	57%	75%

Table IV: Subjects' Choices in the ONE Treatment

Notes: In the ONE Treatment, subjects make each PM and RS choice one single time.

In the Online Appendix, we confirm that the behavior in ONE cannot be rationalized by the behavior in our other treatments. If we take the mixing frequencies observed by each subject in the IND treatment and use that to predict the overall frequency of choices in the ONE treatment, we find that the two are inconsistent: the population of subjects in ONE mixes less than that in IND.

Result 5. *When faced with a single repetition of PM questions, subjects almost always choose the dominant bet. When subjects face RS questions once, they are more likely to choose the risky bet as it becomes more attractive relative to the safe bet. This behavior is not well explained by mixing frequencies in either the SIM or IND treatments.*

Given that subjects choose the dominant bet in the single PM decision but do not choose it in each of twenty repetitions in the IND Experiment, we turn to the SEQ treatment and ask whether it helps subjects view each of the repetitions in isolation, and, as a result, whether it reduces mixing. Moreover, since the SIM block was played

right after the SEQ block, we investigate whether there are spill-over effects between the SEQ and the SIM blocks.

Figure III depicts the percentage of subjects who randomize in the SEQ and SIM treatments and compares these fractions to the IND Experiment.¹⁶ We find that for PM questions, individuals mix less in both the SEQ and the SIM treatments compared with the IND treatment. The reduction in mixing is statistically significant and large in magnitude: the fraction of subjects identified as mixers falls from 64% in the IND experiment to 41% and 45% in the SEQ and SIM treatments, respectively ($p = 0.0019$ IND vs. SEQ, $p = 0.0110$ IND vs. SIM). However, for the RS questions, there are no significant differences in tendency to mix in either the SEQ or SIM treatments compared to the IND experiment ($p = 0.57$ IND vs. SEQ, $p = 0.25$ IND vs. SIM).

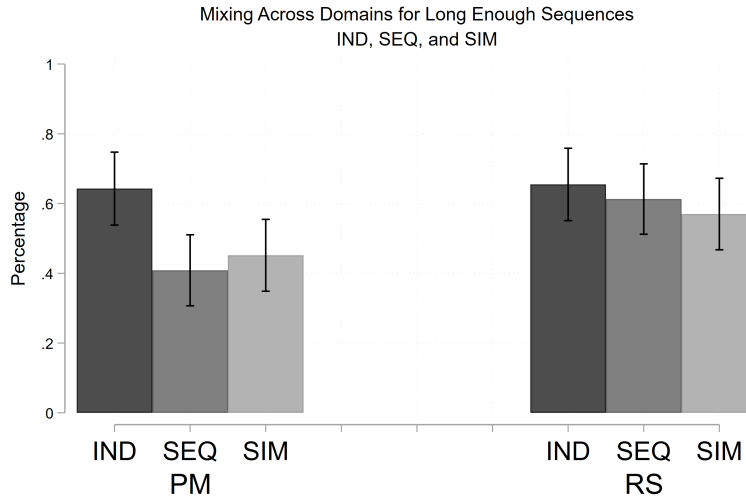


Figure III: Mixing Behavior in IND, SEQ, and SIM Experiments for Long Sequences
Notes: The error bars depict 95% confidence intervals.

¹⁶Given the structure of the SEQ treatment, individuals do not answer all twenty repetitions of a given choice. Therefore, one might worry that the reduction in mixing is an artifact of this “truncation,” where individuals might have mixed had they been given the opportunity to answer more repetitions. To control for this, we identify the position of the average first less-likely bet in the IND treatment. For PM questions, the less-likely bet is always the dominated bet, so we look to see the average first appearance of the dominated bet in a given decision problem. For RS questions, the less-likely bet is the risky bet for low probabilities of the high payoff and is the safe bet for high probabilities of the high payoff. For each question, we identify the less-likely bet and the average first appearance of this bet. We look only at sequences in the SEQ treatment where individuals answered *more* repetitions than this average first less-likely bet. Figure VII in the Online Appendix shows that the results are essentially the same if one looks at the overall data without truncation.

This suggests that the sequential treatment, which eliminates the need for contingent reasoning by design, affects the two types of questions differently—encouraging contingent thinking reduces mixing when mixing is strictly dominated, but it has no effect on mixing when mixing is not dominated.¹⁷ This has implications for identifying the source of mixing in the two types of problems. The significant reduction in mixing in PM-type questions suggests that high mixing frequencies observed in the IND treatment come, at least in part, from the failure of subjects to think about possible contingencies they may face in the future regarding which repetition would be selected for payment. At the same time, since mixing probabilities remain the same in the SEQ, SIM, and IND treatments for RS-type of questions, this suggests that the desire of subjects to randomize in RS-type questions is the manifestation of true underlying preferences. Thus, while both types of mixing were highly prevalent in the IND experiment, and are highly correlated at an individual-level, they seem to stem, in part, from different sources.

Given that individuals mix less in the PM questions, this results in a significantly different type classification of subjects. Since we do not have games in the Sequential Experiment, we re-classify subjects as “Never Mix” if they do not mix in PM or RS questions, “Always Mix” if they mix in both PM and RS, and “Non-Dominated Mix” if they mix only in RS questions. We find significant differences in the type classification of subjects in SEQ and SIM compared to our IND experiment (Fisher’s exact test, IND vs. SEQ $p = 0.032$, IND vs. SIM $p = 0.020$). This results from an increase in *Non-Dominated Mix* subjects replacing a significant portion of the *Always Mix* subjects.

Result 6. *When decisions are sequential, mixing in the PM-type questions is reduced for some, but not all, mixers. Mixing in the RS-type questions, however, is not reduced.*

V.C. Online Experiment

Prelude. One concern is that the mixing we observe is driven by an experimenter demand effect: by asking the same question 20 times, subjects may believe the experimenter *wants* them to mix, and this may drive their behavior. de Quidt et al.

¹⁷See Martínez-Marquina et al. (2019) and literature surveyed there for evidence that people have pervasive difficulties with contingent reasoning. They also find, using a different experimental manipulation, that eliminating the need for contingent reasoning decreases probability matching behavior.

	IND	CORR	SEQ	SIM
Never Mix	29%	20%	27%	35%
Always Mix	58%	56%	46%	38%
Non-Dominated Mix	7%	11%	23%	19%
Others	6%	13%	4%	8%
# of subjects	84	84	93	93

Table V: Individual Types in All Experiments Defined by PM and RS Environments

Notes: We classify subjects as “Never Mix” if they do not mix in PM or RS questions, “Always Mix” if they mix in both PM and RS, and “Non-Dominated Mix” if they mix only in RS questions.

(2018) provide a method for measuring and bounding such experimenter demand effects: run a new treatment with an intentional and clear negative demand effect. If the effect is strong enough then this treatment would provide a lower bound on the true (“demand-free”) level of mixing. It also measures how sensitive mixing behavior is to such demand effects.

Design. This experiment was conducted online because the COVID-19 pandemic had shut down in-person experiment. Thus, we first replicated our IND treatment using the Prolific online platform, and then ran a negative-demand treatment with the same subject pool. We restricted our sample to college-aged subjects who reported United States nationality. Following the suggestion of de Quidt et al. (2018) we added a bold, red sentence to our instructions that read “You would be doing us a favor if you chose the same bet in all 20 choices.” Using MTurk subjects, de Quidt et al. (2018) find that this sentence (modified appropriately for the given task) creates statistically significant shifts in behavior for all 11 tasks they consider, which includes risk, time, ambiguity, and social preferences. The average shift across tasks is 0.6 standard deviations. Thus, if mixing were responsive to experimenter demand effects, we expect a significant drop in mixing behavior compared to our IND replication.

To simplify the experiment, we ran only the PM55, PM80, RS55, and RS80 questions. The order of questions was randomized, as was the location of the bets on the decision screen (left versus right). In total, 82 subjects participated in the replication of the original IND treatment, and 78 participated in the negative-demand treatment.¹⁸

¹⁸We targeted 80 per session based on power calculations for the χ^2 test.




Figure IV: Mixing behavior in an online experiment to test for demand effects.

Results. Figure IV shows the percentage of subjects classified as mixers in the original lab experiment and the two online treatments. It is clear that mixing in the negative-demand treatment (the lightest bars) is not significantly different from the online replication of the IND treatment (the middle bars) in any of the four questions (the χ^2 test p -values are 0.95, 0.66, 0.53, and 0.90, respectively). Following de Quidt et al. (2018), we conclude that mixing behavior does not result from experimenter demand effects. Furthermore, the level of mixing is far from zero: even in PM80 the 95% confidence interval for the negative demand treatment is [19%, 38%]. Thus, in all questions we find that the mixing frequency is significantly positive.

Comparing the online replication to the original lab experiment, it does appear that the mixing propensity is higher on some questions (p -values of 0.597, 0.056, 0.008, and 0.744, respectively), but overall our qualitative results replicate in a very different sample than our original lab environment.¹⁹ Mixing propensities may vary somewhat

¹⁹In a recent working paper, Gupta et al. (2021) find similar results on Prolific as in the University

across populations, but in all cases we find a significantly positive percentage of subjects who mix.

V.D. Implications

As our final result, presented in Table VI, we compare average expected earnings between mixers and non-mixers in all of our treatments. We calculate each individual’s expected payment separately for the PM and RS domain, and we take the average of these expected payments separately for those classified as mixers and non-mixers in that domain.²⁰ We find that mixers have significantly lower expected payments in the PM domain in all treatments. This naturally arises since mixing in PM consists of choosing dominated bets. In contrast, mixers and non-mixers have similar expected payments in the RS domain. This confirms that non-dominated mixing is not objectively detrimental from a payoff perspective, and indeed allows subjects to “smooth” their expected payment.

		IND	CORR	SEQ	SIM	ONLINE	NEG DEMAND
PM	Mixer	\$17.14	\$17.21	\$17.29	\$17.20	\$16.80	\$16.56
	Non-Mixer	\$18.35	\$18.41	\$18.35	\$18.48	\$18.42	\$18.11
	<i>p</i> -value	$p < 0.001$	$p < 0.001$	$p < 0.001$	$p < 0.001$	$p < 0.001$	$p < 0.001$
RS	Mixer	\$16.43	\$16.58	\$16.80	\$16.61	\$16.64	\$16.70
	Non-Mixer	\$16.58	\$16.16	\$16.59	\$16.56	\$16.19	\$16.40
	<i>p</i> -value	$p = 0.52$	$p = 0.091$	$p = 0.30$	$p = 0.84$	$p = 0.079$	$p = 0.28$

Table VI: Average Expected Payments by Domain

Notes: Subjects are classified as a “mixer” or “non-mixer” if they mix in at least one question in the specified domain, irrespective of whether they mix in the other domain.

VI. THEORIES OF RANDOMIZATION

Given that we observe mixing types, an important question is whether there exists a theory that can rationalize such behavior. In this section we review theories of randomization, and show that most cannot explain our data. A more complete and formal analysis of the theories we consider appears in Online Appendix C.

of Pittsburgh lab, so our results are consistent with this.

²⁰Results are indistinguishable if instead we use the classification from Table V and compare Always Mixers to Never Mixers.

VI.A. Models of Preferences Over Reduced Lotteries

Mixing in our experiment generates a two-stage lottery that can be reduced to a simple lottery in the simplex. In this subsection we explore whether observed behavior can be rationalized by a preference relation over these reduced lotteries.

The first important observation is that mixing represents a convex combination of lotteries, so a strict preference for mixing reveals that preferences must be convex. A modest percentage of our subjects never mix, and are therefore consistent with expected utility maximization. But many do mix, and this rules out not only expected utility, but also any model that satisfies the betweenness axiom (Dekel, 1986; Chew, 1989). Examples of models satisfying betweenness include expected utility, weighted utility theory (Chew, 1983), implicit expected utility (Dekel, 1986), skew-symmetric bilinear utility (Fishburn, 1988), Epstein-Zin preferences (Epstein and Zin, 1989), suspicious expected utility (Bordley and Hazen, 1991), and disappointment aversion (Gul, 1991). These models therefore cannot accommodate mixing behavior.

In addition to convexity, rationalizing preferences for subjects who mix in PM questions must also allow for violations of stochastic dominance. This rules out another large class of models, including prospect theory (Kahneman and Tversky, 1979), cumulative prospect theory (Tversky and Kahneman, 1992), rank-dependent expected utility (Quiggin, 1982), quadratic utility (Chew et al., 1991), cautious expected utility (Cerreia-Vioglio et al., 2015), and deliberate randomization (Cerreia-Vioglio et al., 2019). Regret-averse preferences (Loomes and Sugden, 1982), though intransitive, are also ruled out by dominated mixing.

One simple model of convex preferences that does allow for dominance violations is probability weighting, where a (reduced compound) lottery p is evaluated according to $\sum_x w(p(x))u(x)$ for some onto and increasing weighting function $w : [0, 1] \rightarrow [0, 1]$.²¹ If $w(\cdot)$ is chosen appropriately, this model can predict mixing in both PM and RS questions, though we find that the required shape of $w(\cdot)$ needed to fit our data is quite inconsistent with previous estimates of the weighting function; see Online Appendix C for details.

There is, however, an even more fundamental challenge to any rationalizing model

²¹Probability weighting is one component of prospect theory (Kahneman and Tversky, 1979), though prospect theory includes an editing phase that explicitly rules out the choice of dominated lotteries. Models that apply weights to cumulative probabilities—such as rank-dependent utility (Quiggin, 1982) and cumulative prospect theory (Tversky and Kahneman, 1992)—also respect dominance.

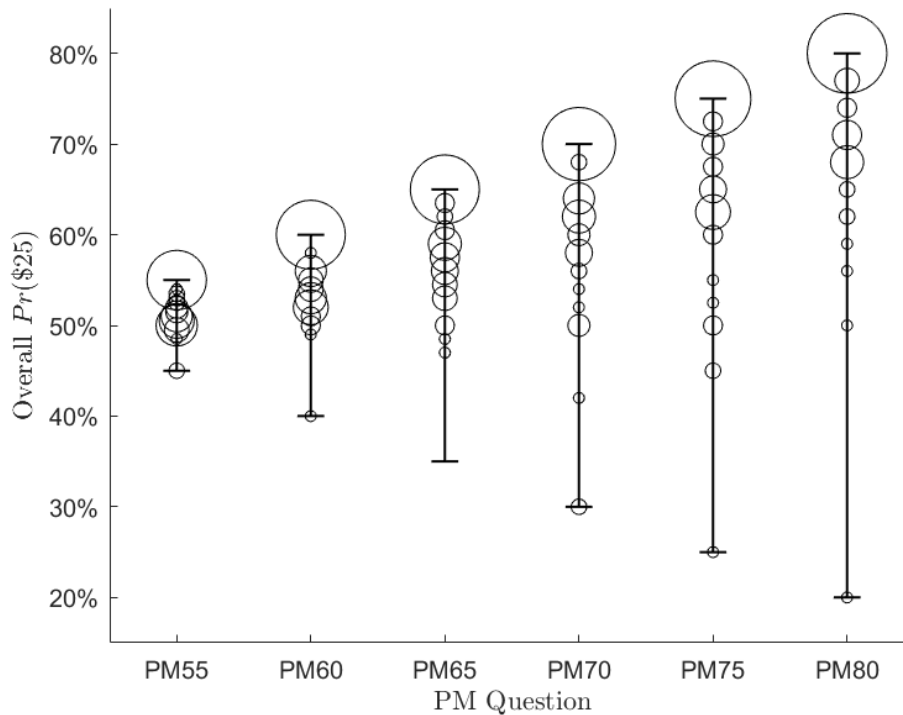


Figure V: Frequencies of overall $Pr(\$25)$ chosen by subjects in each PM question.

Notes: The bubble sizes are proportional to the number of observations.

that assumes reduction, such as the probability weighting theory above: our subjects who mix violate the independence of irrelevant alternatives axiom (IIA, or Property α from Sen, 1969), which is well-known to be necessary for preference maximization. When lotteries are reduced, our PM questions offer menus of lotteries that are nested, which allows us to test IIA. For example, by mixing in PM80 a subject can achieve an overall probability of \$25 anywhere in the range $[0.20, 0.80]$, while in PM75 they can only achieve probabilities of \$25 in the smaller range $[0.25, 0.75]$. These ranges for all six PM questions are shown in Figure V, along with observed choice frequencies. If a subject's choice in PM80 gives a 60% probability of \$25, for example, then IIA requires that their choice in PM75 also gives a 60% probability of \$25. The data shown in Figure V, however, strongly suggests that subjects do not pick the same reduced lottery across nested problems. Subjects who never mix or mix with low frequencies vacuously satisfy IIA, but among those cases where we can test for IIA we find violations in 82% of pairwise comparisons. Thus, the mixing behavior we observe is inconsistent with preference maximization over reduced lotteries.

We also view our data as inconsistent with random utility models for two reasons. First, a random utility model would predict that mixing frequencies are identical between the ONE, IND, and SIM treatments, but we show in Online Appendix B that mixing frequencies are significantly lower in the ONE treatment. Second, as Agranov and Ortoleva (2017) argue, it seems implausible that subjects experience different utility shocks across identical decisions that are made only seconds apart. Our subjects report in the debriefing survey that they consciously randomize, which is at odds with the idea that randomization reflects capricious decision-making.

VI.B. Preferences Over Two-Stage Lotteries

Given that preferences over reduced lotteries cannot explain our data, we now explore whether mixing can be rationalized by a preference over two-stage lotteries that does not respect reduction. In the language of Segal (1990), mixing is a violation of the compound independence axiom. Mixing therefore rules out models that assume compound independence, including expected utility with reduction, recursive expected utility (Kreps and Porteus, 1978; Klibanoff and Ozdenoren, 2007) and models of second-order expected utility (Klibanoff et al., 2005; Ergin and Gul, 2009; Seo, 2009).

One way to accommodate the mixing types we observe is to apply a preference for randomization in the first stage. For example, a perturbed utility model posits that the decision maker chooses a mixture p that maximizes expected utility plus a (typically convex) function $V(p)$ (Machina, 1985; Fudenberg et al., 2015). One interpretation is that $V(p)$ captures the cost of attention or the disutility of effort needed to identify the more-preferred option (Mattsson and Weibull, 2002). An obvious way to apply this concept to two-stage lotteries without reduction is to model the decision maker as having a utility value $U(p)$ for each second-stage lottery p , and choosing the two-stage lottery P that maximizes $\sum_p P(p)U(p) + V(P)$. This is exactly the approach of Allen and Rehbeck (2019), who make no assumptions on $U(p)$. Siegel (1961, pp. 769–770) proposes a specific perturbed utility model in which $U(p)$ is simply the expected value of p and $V(P)$ rewards variance in P .

Allen and Rehbeck (2019) provide a revealed-preference test of perturbed utility models of this form that applies easily to our RS questions. Recall that the risky alternative in RS75 pays \$25 for balls 1–15, while the risky alternative in RS80 pays \$25 for balls 1–16. If we assume $U(\cdot)$ gives a higher value to the latter, then the

Allen and Rehbeck (2019) condition is that the subject must pick the risky alternative more frequently in RS80 than in RS75. Indeed, this must be true for any pair of RS questions, and there are 15 possible such comparisons. Subjects who never mix satisfy this condition in all 15 comparisons. Of the 65% whom we classify as mixers in the RS domain, however, only 18% satisfy the condition in all 15 comparisons. But, among all mixers, the average number of failures per person (out of 15) is only 2.67, and the vast majority of those are in adjacent RS questions (for example, RS75 vs. RS80).²² Thus, while we don't see perfect support, a model of preferences for first-stage randomization without reduction fits the behavior of our mixing types reasonably well.

Although such models can organize our mixing data, they are perhaps a bit unsatisfying in that they do not provide an underlying justification for mixing. Instead, they simply posit that mixing is chosen because the subject has a utility for mixing. Whether one could derive this utility from some deeper motivation remains an interesting open question.

VI.C. Mistakes, Biases, and Heuristics

While a preference for first-stage randomization can explain many of those who mix in all situations, it seems unable to explain the type of subject who mixes in IND but not in SEQ. For these subjects, mixing appears to be a mistake or heuristic that is overturned when decisions are made sequentially. Here we review a variety of plausible heuristics and biases that can lead to randomization behavior.

One well-known bias that can predict mixing is the gambler's fallacy: subjects wrongly believe that draws from the bingo cage exhibit negative serial correlation (Rabin and Vayanos, 2010, e.g.). For example, a subject who has chosen the dominant bet in a PM problem a few times in a row—and believes those bets are likely to pay off—might think that the dominated bet is now “due” to pay off, leading them to switch the dominated bet. Indeed, we show in Online Appendix C that if the belief in negative serial correlation is high enough then the optimal strategy in the IND treatment is to alternate between bets across replications. In the CORR treatment, however, there is only one draw from the bingo cage, so there is no way for negative correlation to affect betting behavior: a subject who believes the dominant bet is the

²²Results in the SEQ treatment are similar: Of the 68% who are mixers (classified as such regardless of the number of replicates faced), only 16% satisfy the condition on all 15 comparisons, but the average number of failures per person is only 3.44.

better bet must believe this on all twenty replications. Yet we still find significant mixing behavior, so a simple gambler's fallacy story cannot explain our data.

Rabin (2002) describes a model in which subjects wrongly believe draws are made with replacement, which introduces a different kind of serial correlation. In the appendix we show that this also cannot explain the mixing we observe.

Motivated by the law of small numbers (Tversky and Kahneman, 1971; Rabin, 2002) we also consider a theory which we call the "modal count heuristic." According to this theory, a subject in PM60, for example, correctly identifies that the modal number of times the dominant bet will pay off is twelve out of twenty. Based on this, they choose the dominant bet twelve times. Their mistake is in failing to realize that they are unlikely to predict *which* twelve replicates are the ones that will pay off. While this mistake can explain mixing in the IND treatment, it cannot explain mixing in the CORR treatment because there the dominant bet either pays off twenty times or zero times.

Another theory (which we develop more completely in the appendix) is one of regret with a convex cost of "mistakes." Consider a PM decision problem in the CORR treatment in which the subject chooses the dominant bet in all twenty replications. If the realized draw is such that the dominated bet was the better bet then, *ex-post*, this subject has made twenty "mistakes." If the subject has a convex cost of such mistakes, and *ex-ante* makes decisions accounting for their expected cost of mistakes, then their optimal strategy may not be to choose the dominant bet in all twenty replications. Instead, they may prefer to mix in the dominated bet in order to reduce the number of mistakes they would make in each state of the world. Thus, this theory predicts mixing in the CORR treatment. In the IND treatment, however, every time the subject chooses the dominated bet they increase the probability of making a mistake on that replicate, without affecting their probabilities of mistakes on other replicates. Since the expected cost of mistakes always increases by choosing the dominated bet, this theory cannot explain mixing in the IND treatment.

Another possibility is that subjects exhibit irrational diversification (Read and Loewenstein, 1995; Baltussen and Post, 2011; Rubinstein, 2002). For example, the subject may incorrectly believe that they are paid for all twenty choices instead of one randomly-selected choice. In the CORR treatment, choosing different bets on a PM question allows the subject to hedge against the single realization of uncertainty. In the IND treatment, however, each bet's payoff is determined by an independent

realization of uncertainty, so there is no opportunity to hedge. Thus, if irrational diversification does explain mixing then it must be purely irrational, stemming from a direct preference for mixing rather than from a false perception that mixing somehow improves the lottery they receive.

The fact that some types of subjects mix in IND but not in SEQ is reminiscent of the observation that subjects overbid in the simultaneous-move second-price auction but bid truthfully in the dynamic English auction. Li (2017) argues this is because truth-telling is obviously dominant in the latter, but not the former. Roughly, this is because the worst-case outcome under truthful bidding is preferred to the best-case outcome under any deviation. Unfortunately, this logic does not apply to mixing in our PM questions. If we view the outcome as “which replicate is chosen for payment,” then choosing the more-preferred option all twenty times is obviously dominant in both IND and SEQ. If instead we view the outcome as “which replicate is chosen *and* which ball is drawn” then no vector of choices obviously dominates another in either the IND or SEQ treatments. See the appendix for details.

Unfortunately, none of these heuristics or biases explain both mixing in IND and the reduction of mixing in SEQ. A simple explanation is that subjects have a preference for randomization in the IND treatment because of decision costs or inattention (captured by the perturbed utility model discussed above), and the SEQ treatment helps reduce these costs by focusing the subject on one question at a time.

Some authors argue that probability matching and randomization behavior have evolutionary foundations (Cooper, 1989; Gigerenzer, 2002; Brennan and Lo, 2012). Probability matching can be thought of as an evolutionarily stable strategy (Fretwell, 1972), leading to its persistence in decision making. These models assume, however, that risks are perfectly correlated across agents: for example, if rational agents pick the dominant option all 20 times in a PM question and then happen to get a bad draw then all rational agents are wiped out and do not produce offspring. But if risks are independent across agents then the rational types would be most likely to thrive.

Regardless, our experiment cannot test evolutionary explanations; if there is an evolutionary mechanism behind randomization then it would simply show up as an underlying preference for randomization, which is exactly what we see in some agents. But any evolutionary explanation would also need to account for the heterogeneity we observe: we find that mixers and non-mixers coexist in the population.

Finally, we are unaware of any theory that explicitly predicts the increase in mixing

we observe in games, compared to individual decision problems. One conjecture is that games introduce ambiguity (in the form of strategic uncertainty), and ambiguity induces randomization behavior. This might be similar to the discussions of “source functions” and preferences dependent on the source of uncertainty (Tversky and Fox, 1995; Abdellaoui et al., 2011). We discuss this in more detail in our theoretical appendix. We view this as a promising direction for future research.

VII. DISCUSSION AND CONCLUDING THOUGHTS

We study individuals’ tendency to randomize their choices, documenting patterns within domains and correlations across domains. Randomization is ubiquitous, but systematic. Randomization is highly correlated within individual, responds monotonically to parameter changes in the environment, and increases in strategic situations. Few theories in the literature can accommodate our results, other than those that directly assume a preference for mixing and little else. When choices are sequential we find that for some subjects mixing is reduced in decisions with a dominated option, but not in risky-safe decisions. This effect persists when decisions are then made simultaneously, suggesting that dominated mixing was more of a heuristic for these individuals, rather than a preference.

Our results highlight a number of open questions. First, it would be interesting to elicit subjects’ “strategies” for making these decisions. This would allow us to see whether individuals believe they “should” randomize in these environments. Similarly, it would be interesting to understand whether randomization is normative; for example, would individuals choose to randomize for others? Second, we provide correlations between risk preferences and randomization in the appendix, but it would be interesting to learn more about the relationship between risk preferences and diversification behavior in these environments. One could imagine risk-averse individuals randomizing in order to hedge, but one could also imagine risk-seeking individuals randomizing to increase risk. Finally, it would be interesting to study other interventions that reduce randomization and identify conditions under which risky-safe randomization also disappears, if such conditions exist.

The results from our SEQ treatment contribute to a growing experimental literature that investigates interventions that reduce violations of dominance (Charness et al., 2007; Schulze and Newell, 2016) and, more generally, explore subjects’ ability to

reason about future or hypothetical contingencies (Li, 2017; Esponda and Vespa, 2019; Martínez-Marquina et al., 2019). Esponda and Vespa (2019) show that violations of Savage’s sure-thing principle are reduced when subjects are primed to ignore the “irrelevant” states in a given decision. This is similar to our SEQ treatment which encourages subjects to focus only on the current replicate. Charness et al. (2007) find that violations of dominance are reduced when Bayesian updating is not required. Probability matching has been shown to reduce when subjects are first primed to think about payoff-maximizing strategies (Gal and Baron, 1996; Newell et al., 2013; Koehler and James, 2010) or are asked to recommend a strategy to another player (Fantino and Esfandiari, 2002).

The success of these interventions suggests that probability matching is likely a mistake.²³ Our experiments corroborate this view and add additional insights. Our SEQ treatment identifies the crucial difference in mechanisms driving behavior in choices that involve dominated lotteries and those that do not. We show that it is possible to “train away” randomization for at least a third of the subjects, but only when choices involve dominated lotteries. This suggests that mixing between risky and safe lotteries may not stem from a failure of contingent reasoning, but may indicate the true desire to choose both alternatives with positive probability.

²³See also Nielsen and Rehbeck (2020), who elicit which axioms subjects wish to follow and, when their subsequent choices violate the axiom, whether they wish to change either of these decisions. They include a “consistency” axiom, which implies making the same choice in two identical decisions, and ask the same lottery choice twice randomly within a set of 33 lottery choices. They find that consistency violations are usually viewed as a mistake, with subjects changing their lottery choices to make the same choice in both decisions.

REFERENCES

- ABDELLAOUI, M., A. BAILLON, L. PLACIDO, AND P. P. WAKKER (2011): “The Rich Domain of Uncertainty: Source Functions and Their Experimental Implementation,” *American Economic Review*, 101, 695–723.
- ACKERT, L. F., A. B. GILLETTE, J. MARTINEZ-VAZQUEZ, AND M. RIDER (2011): “Are Benevolent Dictators Altruistic in Groups? A within-Subject Design,” *Experimental Economics*, 14, 307–321.
- AGRANOV, M. AND P. ORTOLEVA (2017): “Stochastic Choice and Preferences for Randomization,” *Journal of Political Economy*, 125, 40–68.
- (2022): “Ranges of Preferences and Randomization,” Caltech working paper.
- ALLEN, R. AND J. REHBECK (2019): “Revealed Stochastic Choice with Attributes,” Ohio State University working paper.
- (2021): “A Generalization of Quantal Response Equilibrium via Perturbed Utility,” *Games*, 12, 20.
- ANDERSON, L. R. AND J. M. MELLOR (2009): “Are Risk Preferences Stable? Comparing an Experimental Measure with a Validated Survey-Based Measure,” *Journal of Risk and Uncertainty*, 39, 137–160.
- AZRIELI, Y., C. P. CHAMBERS, AND P. J. HEALY (2018): “Incentives in Experiments: A Theoretical Analysis,” *Journal of Political Economy*, 126, 1472–1503.
- BALTUSSEN, G. AND G. T. POST (2011): “Irrational Diversification: An Examination of Individual Portfolio Choice,” *Journal of Financial and Quantitative Analysis*, 46, 1463–1491.
- BARTLING, B. AND U. FISCHBACHER (2012): “Shifting the Blame: On Delegation and Responsibility,” *The Review of Economic Studies*, 79, 67–87.
- BECKER, G. S. (1968): “Crime and Punishment: An Economic Approach,” *Journal of Political Economy*, 76, 169.

- BERG, J., J. DICKHAUT, AND K. MCCABE (2005): “Risk Preference Instability across Institutions: A Dilemma,” *Proceedings of the National Academy of Sciences*, 102, 4209–4214.
- BINSWANGER, H. P. (1980): “Attitudes toward Risk: Experimental Measurement in Rural India,” *American journal of agricultural economics*, 62, 395–407.
- BLANCO, M., D. ENGELMANN, AND H. T. NORMANN (2011): “A Within-Subject Analysis of Other-Regarding Preferences,” *Games and Economic Behavior*, 72, 321–338.
- BORDLEY, R. AND G. B. HAZEN (1991): “SSB and Weighted Linear Utility as Expected Utility with Suspicion,” *Management Science*, 37, 396–408.
- BRENNAN, T. J. AND A. W. LO (2012): “An Evolutionary Model of Bounded Rationality and Intelligence,” *PLOS One*.
- BROSIG, J., T. REICHMANN, AND J. WEIMANN (2007): “Selfish in the End?” *Working Paper Series*.
- BROWN, A. L. AND P. J. HEALY (2018): “Separated Decisions,” *European Economic Review*, 101, 20–34.
- BRUHIN, A., E. FEHR, AND D. SCHUNK (2019): “The Many Faces of Human Sociality: Uncovering the Distribution and Stability of Social Preferences,” *Journal of the European Economic Association*, 17, 1025–1069.
- CAMERER, C. F. AND T.-H. HO (1994): “Violations of the Betweenness Axiom and Nonlinearity in Probability,” *Journal of risk and uncertainty*, 8, 167–196.
- CERREIA-VIOGLIO, S., D. DILLENBERGER, AND P. ORTOLEVA (2015): “Cautious Expected Utility and the Certainty Effect,” *Econometrica*, 83, 693–728.
- CERREIA-VIOGLIO, S., D. DILLENBERGER, P. ORTOLEVA, AND G. RIELLA (2019): “Deliberately Stochastic,” *American Economic Review*, 109, 2425–45.
- CHARNESS, G., U. GNEEZY, AND A. IMAS (2013): “Experimental Methods: Eliciting Risk Preferences,” *Journal of Economic Behavior and Organization*, 87, 43–51.

- CHARNESS, G., E. KARNI, AND D. LEVIN (2007): “Individual and Group Decision Making under Risk: An Experimental Study of Bayesian Updating and Violations of First-Order Stochastic Dominance,” *Journal of Risk and uncertainty*, 35, 129–148.
- CHEW, S.-H. (1983): “A Generalization of the Quasilinear Mean with Applications to the Measurement of Income Inequality and Decision Theory Resolving the Allais Paradox,” *Econometrica*, 1065–1092.
- CHEW, S. H. (1989): “Axiomatic Utility Theories with the Betweenness Property,” *Annals of operations Research*, 19, 273–298.
- CHEW, S. H., L. G. EPSTEIN, AND U. SEGAL (1991): “Mixture Symmetry and Quadratic Utility,” *Econometrica*, 59, 139–163.
- CHOI, S., R. FISMAN, D. GALE, AND S. KARIV (2007): “Consistency and Heterogeneity of Individual Behavior under Uncertainty,” *American economic review*, 97, 1921–1938.
- CHUANG, Y. AND L. SCHECHTER (2015): “Stability of Experimental and Survey Measures of Risk, Time, and Social Preferences: A Review and Some New Results,” *Journal of Development Economics*, 117, 151–170.
- COOPER, W. S. (1989): “How Evolutionary Biology Challenges the Classical Theory of Rational Choice,” *Biology and Philosophy*, 4, 457–481.
- CROSETTO, P. AND A. FILIPPIN (2016): “A Theoretical and Experimental Appraisal of Four Risk Elicitation Methods,” *Experimental Economics*, 19, 613–641.
- DAVE, C., C. C. ECKEL, C. A. JOHNSON, AND C. ROJAS (2010): “Eliciting Risk Preferences: When Is Simple Better?” *Journal of Risk and Uncertainty*, 41, 219–243.
- DE OLIVEIRA, A., R. T. CROSON, AND C. C. ECKEL (2012a): “Are Preferences Stable across Domains? An Experimental Investigation of Social Preferences in the Field,” *Southern Economic Journal*, 79.
- DE OLIVEIRA, A. C., C. ECKEL, AND R. T. CROSON (2012b): “The Stability of Social Preferences in a Low-Income Neighborhood,” *Southern Economic Journal*, 79, 15–45.

- DE QUIDT, J., J. HAUSHOFER, AND C. ROTH (2018): “Measuring and Bounding Experimenter Demand,” *American Economic Review*, 108, 3266–3302.
- DEKEL, E. (1986): “An Axiomatic Characterization of Preferences under Uncertainty: Weakening the Independence Axiom,” *Journal of Economic Theory*, 40, 304–318.
- DIECIDUE, E., U. SCHMIDT, AND P. P. WAKKER (2004): “The Utility of Gambling Reconsidered,” *Journal of Risk and Uncertainty*, 29, 241–259.
- DULLECK, U., J. FOOKEN, AND J. FELL (2015): “Within-Subject Intra- and Inter-Method Consistency of Two Experimental Risk Attitude Elicitation Methods,” *German Economic Review*, 16, 104–121.
- DWENGER, N., D. KÜBLER, AND G. WEIZSÄCKER (2018): “Flipping a Coin: Evidence from University Applications,” *Journal of Public Economics*, 167, 240–250.
- ECKEL, C. C. AND R. K. WILSON (2004): “Is Trust a Risky Decision?” *Journal of Economic Behavior & Organization*, 55, 447–465.
- EPSTEIN, L. G. AND S. E. ZIN (1989): “Substitution, Risk Aversion, and the Temporal Behavior of Consumption: A Theoretical Framework,” *Econometrica*, 57, 937–969.
- ERGIN, H. AND F. GUL (2009): “A Theory of Subjective Compound Lotteries,” *Journal of Economic Theory*, 144, 899–929.
- ESPONDA, I. AND E. VESPA (2019): “Contingent Thinking and the Sure-Thing Principle: Revisiting Classic Anomalies in the Laboratory,” UC Santa Barbara working paper.
- FANTINO, E. AND A. ESFANDIARI (2002): “Probability Matching: Encouraging Optimal Responding in Humans,” *Canadian Journal of Experimental Psychology*, 56, 58.
- FELDMAN, P. AND J. REHBECK (2019): “Revealing a Preference for Mixing: An Experimental Study of Risk,” Ohio State University working paper.
- FISHBURN, P. C. (1988): *Nonlinear Preference and Utility Theory*, 5, Johns Hopkins University Press Baltimore.

- FISMAN, R., S. KARIV, AND D. MARKOVITS (2007): “Individual Preferences for Giving,” *American Economic Review*, 97, 1858–1876.
- FRETWELL, S. D. (1972): *Populations in Seasonal Environments*, Princeton University Press.
- FRIEDMAN, E. AND J. WARD (2022): “Stochastic Choice and Noisy Beliefs in Games: An Experiment,” University of Essex working paper.
- FUDENBERG, D., R. IIJIMA, AND T. STRZALECKI (2015): “Stochastic Choice and Revealed Perturbed Utility,” *Econometrica*, 83, 2371–2409.
- GAL, I. AND J. BARON (1996): “Understanding Repeated Simple Choices,” *Thinking & Reasoning*, 2, 81–98.
- GEORGANAS, S., P. J. HEALY, AND R. A. WEBER (2015): “On the Persistence of Strategic Sophistication,” *Journal of Economic Theory*, 159, 369–400.
- GIGERENZER, G. (2002): *Adaptive Thinking: Rationality in the Real World*, Oxford University Press.
- GILLEN, B., E. SNOWBERG, AND L. YARIV (2019): “Experimenting with Measurement Error: Techniques with Applications to the Caltech Cohort Study,” *Journal of Political Economics*, 127, 1826–1863.
- GNEEZY, U. AND J. POTTERS (1997): “An Experiment on Risk Taking and Evaluation Periods,” *The Quarterly Journal of Economics*, 112, 631–645.
- GOEREE, J. K., C. A. HOLT, AND T. R. PALFREY (2005): “Regular Quantal Response Equilibrium,” *Experimental Economics*, 8, 347–367.
- GRANT, D., H. HAKE, AND J. HORNSETH (1951): “Acquisition and Extinction of Verbal Expectations in a Situation Analogous to Conditioning,” *Journal of Experimental Psychology*, 42, 5.
- GREINER, B. (2015): “Subject Pool Recruitment Procedures: Organizing Experiments with ORSEE,” *Journal of the Economic Science Association*, 1, 114–125.
- GRETHER, D. M. (1981): “Financial Incentive Effects and Individual Decisionmaking,” Caltech working paper.

- GUL, F. (1991): “A Theory of Disappointment Aversion,” *Econometrica: Journal of the Econometric Society*, 667–686.
- GUPTA, N., L. RIGOTTI, AND A. WILSON (2021): “The Experimenters’ Dilemma: Inferential Preferences over Populations,” *arXiv preprint arXiv:2107.05064*.
- HALEVY, Y. (2015): “Time Consistency: Stationarity and Time Invariance,” *Econometrica*, 83, 335–352.
- HARRISON, G. W., M. LAU, E. RUTSTROM, AND M. B. SULLIVAN (2005): “Eliciting Risk and Time Preferences Using Field Experiments: Some Methodological Issues,” in *Field Experiments in Economics*, ed. by J. Carpenter, G. Harrison, and J. List, Emerald Group Publishing Limited, vol. 10 of *Research in Experimental Economics*.
- HEY, J. (2001): “Does Repetition Improve Consistency?” *Experimental Economics*, 4, 5–54.
- HOROWITZ, J. K. (1992): “A Test of Intertemporal Consistency,” *Journal of Economic Behavior & Organization*, 17, 171–182.
- HUMPHREYS, L. G. (1939): “Acquisition and Extinction of Verbal Expectations in a Situation Analogous to Conditioning,” *Journal of Experimental Psychology*, 25, 294.
- ISAAC, R. M. AND D. JAMES (2000): “Just Who Are You Calling Risk Averse?” *Journal of Risk and Uncertainty*, 20, 177–187.
- KAHNEMAN, D. AND A. TVERSKY (1979): “Prospect Theory: An Analysis of Decision under Risk,” *Econometrica*, 47, 263–291.
- KIRBY, K. N. (2009): “One-Year Temporal Stability of Delay-Discount Rates,” *Psychonomic Bulletin & Review*, 16, 457–462.
- KLIBANOFF, P., M. MARINACCI, AND S. MUKERJI (2005): “A Smooth Model of Decision Making under Ambiguity,” *Econometrica*, 73, 1849–1892.
- KLIBANOFF, P. AND E. OZDENOREN (2007): “Subjective Recursive Expected Utility,” *Economic Theory*, 30, 49–87.
- KOEHLER, D. J. AND G. JAMES (2010): “Probability Matching and Strategy Availability,” *Memory & cognition*, 38, 667–676.

- KREPS, D. M. AND E. L. PORTEUS (1978): “Temporal Resolution of Uncertainty and Dynamic Choice Theory,” *Econometrica*, 46, 185–200.
- KRUPKA, E. L. AND M. STEPHENS (2013): “The Stability of Measured Time Preferences,” *Journal of Economic Behavior & Organization*, 85, 11–19.
- KRUSE, J. B. AND M. A. THOMPSON (2003): “Valuing Low Probability Risk: Survey and Experimental Evidence,” *Journal of Economic Behavior & Organization*, 50, 495–505.
- LEVITT, S. D. (2021): “Heads or Tails: The Impact of a Coin Toss on Major Life Decisions and Subsequent Happiness,” *The Review of Economic Studies*, 88, 378–405.
- LI, S. (2017): “Obviously Strategy-Proof Mechanisms,” *American Economic Review*, 107, 3257–87.
- LÖNNQVIST, J.-E., M. VERKASALO, G. WALKOWITZ, AND P. C. WICHARDT (2015): “Measuring Individual Risk Attitudes in the Lab: Task or Ask? An Empirical Comparison,” *Journal of Economic Behavior & Organization*, 119, 254–266.
- LOOMES, G. (1998): “Probabilities vs Money: A Test of Some Fundamental Assumptions about Rational Decision Making,” *Economic Journal*, 108, 477–489.
- LOOMES, G. AND R. SUGDEN (1982): “Regret Theory: An Alternative Theory of Rational Choice Under Uncertainty,” *Economic Journal*, 92, 805–824.
- LOTZ, S., T. SCHLÖSSER, D. M. CAIN, AND D. FETCHENHAUER (2013): “The (in)Stability of Social Preferences: Using Justice Sensitivity to Predict When Altruism Collapses,” *Journal of Economic Behavior & Organization*, 93, 141–148.
- MACHINA, M. J. (1985): “Stochastic Choice Functions Generated from Deterministic Preferences over Lotteries,” *The economic journal*, 95, 575–594.
- MARTÍNEZ-MARQUINA, A., M. NIEDERLE, AND E. VESPA (2019): “Failures in Contingent Reasoning: The Role of Uncertainty,” *American Economic Review*, 109, 3437–74.
- MATTSSON, L.-G. AND J. W. WEIBULL (2002): “Probabilistic Choice and Procedurally Bounded Rationality,” *Games and Economic Behavior*, 41, 61–78.

- MCKELVEY, R. D. AND T. R. PALFREY (1995): “Quantal Response Equilibria for Normal Form Games,” *Games and Economic Behavior*, 10, 6–38.
- MCLEISH, K. N. AND R. J. OXOBY (2007): “Gender, Affect and Intertemporal Consistency: An Experimental Approach,” SSRN Scholarly Paper ID 977508, Social Science Research Network, Rochester, NY.
- MEIER, S. AND C. D. SPRENGER (2014): “Temporal Stability of Time Preferences,” *The Review of Economics and Statistics*, 97, 273–286.
- NEILSON, W. S. (1992): “Some Mixed Results on Boundary Effects,” *Economics Letters*, 39, 275–278.
- NEWELL, B. R., D. J. KOEHLER, G. JAMES, T. RAKOW, AND D. VAN RAVENZWAAIJ (2013): “Probability Matching in Risky Choice: The Interplay of Feedback and Strategy Availability,” *Memory & Cognition*, 41, 329–338.
- NIELSEN, K. AND J. REHBECK (2020): “When Choices Are Mistakes,” Ohio State University working paper.
- QUIGGIN, J. (1982): “A Theory of Anticipated Utility,” *Journal of Economic Behavior and Organization*, 3, 323–343.
- RABIN, M. (2002): “Inference by Believers in the Law of Small Numbers,” *The Quarterly Journal of Economics*, 117, 775–816.
- RABIN, M. AND D. VAYANOS (2010): “The Gambler’s and Hot-Hand Fallacies: Theory and Applications,” *The Review of Economic Studies*, 77, 730–778.
- READ, D. AND G. LOEWENSTEIN (1995): “Diversification Bias: Explaining the Discrepancy in Variety Seeking between Combined and Separated Choices,” *Journal of Experimental Psychology: Applied*, 1, 34.
- REUBEN, E., P. SAPIENZA, AND L. ZINGALES (2010): “Time Discounting for Primary and Monetary Rewards,” *Economics Letters*, 106, 125–127.
- REYNAUD, A. AND S. COUTURE (2012): “Stability of Risk Preference Measures: Results from a Field Experiment on French Farmers,” *Theory and Decision*, 73, 203–221.

- ROMERO, J. AND Y. ROSOKHA (2019): “The Evolution of Cooperation: The Role of Costly Strategy Adjustments,” *American Economic Journal: Microeconomics*, 11, 299–328.
- RUBINSTEIN, A. (2002): “Irrational Diversification in Multiple Decision Problems,” *European Economic Review*, 46, 1369–1378.
- SANDRONI, A., S. LUDWIG, AND P. KIRCHER (2013): “On the Difference between Social and Private Goods,” *The B.E. Journal of Theoretical Economics*, 13, 151–177.
- SCHMIDT, U. (1998): “A Measurement of the Certainty Effect,” *Journal of Mathematical Psychology*, 42, 32–47.
- SCHULZE, C. AND B. R. NEWELL (2016): “More Heads Choose Better than One: Group Decision Making Can Eliminate Probability Matching,” *Psychonomic Bulletin & Review*, 23, 907–914.
- SEGAL, U. (1990): “Two-Stage Lotteries without the Reduction Axiom,” *Econometrica*, 349–377.
- SEN, A. (1969): “Quasi-Transitivity, Rational Choice and Collective Decisions,” *The Review of Economic Studies*, 36, 381–393.
- SEO, K. (2009): “Ambiguity and Second-Order Belief,” *Econometrica*, 77, 1575–1605.
- SHACHAT, J. M. (2002): “Mixed Strategy Play and the Minimax Hypothesis,” *Journal of Economic Theory*, 104, 189–226.
- SIEGEL, S. (1961): “Decision Making and Learning under Varying Conditions of Reinforcement.” *Annals of the New York Academy of Sciences*.
- SIEGEL, S. AND D. A. GOLDSTEIN (1959): “Decision-Making Behavior in a Two-Choice Uncertain Outcome Situation.” *Journal of Experimental Psychology*, 57, 37.
- SLOVIC, P. (1972): “Information Processing, Situation Specificity, and the Generality of Risk-Taking Behavior,” *Journal of Personality and Social Psychology*, 22, 128–134.

- SOPHER, B. AND M. NARRAMORE (2000): “Stochastic Choice and Consistency in Decision Making under Risk: An Experimental Study,” *Theory and Decision. An International Journal for Multidisciplinary Advances in Decision Science*, 48, 323–349.
- TVERSKY, A. AND C. R. FOX (1995): “Weighing Risk and Uncertainty.” *Psychological review*, 102, 269.
- TVERSKY, A. AND D. KAHNEMAN (1971): “Belief in the Law of Small Numbers.” *Psychological bulletin*, 76, 105.
- (1992): “Advances in Prospect Theory: Cumulative Representation of Uncertainty,” *Journal of Risk and uncertainty*, 5, 297–323.
- UBFAL, D. (2016): “How General Are Time Preferences? Eliciting Good-Specific Discount Rates,” *Journal of Development Economics*, 118, 150–170.
- VLAEV, I., N. CHATER 1, AND N. STEWART (2009): “Dimensionality of Risk Perception: Factors Affecting Consumer Understanding and Evaluation of Financial Risk,” *Journal of Behavioral Finance*, 10, 158–181.
- WU, G. AND R. GONZALEZ (1996): “Curvature of the Probability Weighting Function,” *Management Science*, 42, 1676–1690.
- ZHANG, X. AND S. ZHONG (2020): “Putting Preference for Randomization to Work,” National University of Singapore working paper.

For Online Publication Only

A. BELIEF PAYMENT PROCEDURE

To elicit subjects' beliefs, we ask them to imagine filling out a table like the one shown below. In each row, a subject chooses Option A (which pays if a randomly-selected column player chooses Left) or Option B (which pays with the given probability, as determined by rolls of dice). Rather than eliciting all 100 responses, we assume subjects would start out preferring Option A and at some point would switch to choosing Option B. We ask subjects to report the row—or probability of receiving \$20—at which they would switch from choosing Option A to choosing Option B. This is the subject's belief that Column will choose Left. Call this belief p .

Q#		Option A		Option B
1	Would you rather have	\$20 if Column chose Left	or	1% chance of \$20
2	Would you rather have	\$20 if Column chose Left	or	2% chance of \$20
3	Would you rather have	\$20 if Column chose Left	or	3% chance of \$20
⋮	⋮	⋮	⋮	⋮
99	Would you rather have	\$20 if Column chose Left	or	99% chance of \$20
100	Would you rather have	\$20 if Column chose Left	or	100% chance of \$20

Table VII: Belief Elicitation Questions

If this belief elicitation were chosen for payment, we would use dice to draw a uniform random number $x \in \{1, \dots, 100\}$ and pay based on what the subject chose in row x . Thus, if $x < p$ the subject would receive Option A: \$20 if a randomly-selected column player chose Left. If $x \geq p$ then the subject would receive Option B: \$20 with probability $x\%$. This lottery is also resolved using die rolls: We use dice to draw a number uniformly from $\{1, \dots, 100\}$ and pay \$20 if the number drawn is less than x .

B. ADDITIONAL RESULTS

B.1. Mixing in ONE Compared to IND and SIM

We verify here that choice frequencies in the ONE treatment cannot be rationalized as being consistent with randomization behavior in either the SIM or IND treatment. Consider first the comparison with the SIM treatment, which uses the same set of subjects. Let p_{ij}^{SIM} be the probability with which subject i chooses the dominant or

risky bet on decision problem j in the SIM treatment. For simplicity, we assume this is perfectly measured by the fraction of the 20 replicates in which the subject chose the dominant or risky bet.²⁴ In the ONE treatment, let $x_{ij} = 1$ if i chooses the dominant or risky bet, and $x_{ij} = 0$ otherwise. Under the null hypothesis that subjects randomize equally in both treatments, $x_j = \sum_i x_{ij}$ is distributed according to a Poisson binomial distribution with mean $\mu_j = \sum_i p_{ij}^{SIM}$ and variance $\sigma_j^2 = \sum_i p_{ij}^{SIM}(1 - p_{ij}^{SIM})$. This is well-approximated by a normal distribution with mean μ_j and variance σ_j^2 , whose cdf we denote by $\Phi(\cdot|\mu_j, \sigma_j)$. Thus, we can reject this null at the 5% level for each question j if $\Phi(x_j|\mu_j, \sigma_j) < 0.025$ or $\Phi(x_j|\mu_j, \sigma_j) > 0.975$.

		Probability (%)					
		55%	60%	65%	70%	75%	80%
ONE vs. SIM	PM	0.998***	0.987**	0.998***	0.999***	0.990***	0.975**
	RS	0.028*	0.157	0.006***	0.980**	0.923	>0.999***
ONE vs. IND	PM	>0.999***	0.999***	>0.999***	>0.999***	>0.999***	0.986***
	RS	0.013*	0.202	0.146	0.967	0.918	>0.999***

Table VIII: The probability of observing choice frequencies equal to or less than those observed in the ONE treatment under the null hypothesis that mixing frequencies are equal across treatments. Significance levels (two-tailed test): *10%, **5%, ***1%.

In the first row of Table VIII we see that $\Phi(x_j|\mu_j, \sigma_j) \geq 0.975$ for all six PM questions, leading us to conclude that behavior in ONE cannot be explained by the mixing behavior observed in SIM for any PM questions. In particular, subjects in the ONE treatment choose the dominant option far more frequently than implied by mixing frequencies in the SIM treatment. For the RS questions, however, results vary by question. The risky choice is chosen less often than predicted in the first three questions, and more often in the last three questions, though statistical significance varies.²⁵

Comparing the ONE treatment to the IND treatment is more difficult because the sets of subjects differ between these treatments. Indeed, the IND treatment has 84

²⁴See footnote 25 for a robustness check of this assumption.

²⁵Here we assumed p_{ij}^{SIM} is perfectly measured from behavior in the SIM treatment. An alternative model is that each p_{ij}^{SIM} is uniformly distributed, but upon observing subject i 's behavior in SIM we form a Bayesian posterior about p_{ij}^{SIM} . Monte Carlo simulations show that we reject this null hypotheses even more often than in the perfectly-measured hypothesis. This is because, under the Bayesian model, each μ_j is pushed closer to the prior mean of $\sum_i(1/2)$, while the actual ONE data lie in the opposite direction.

subjects while the ONE treatment has 93. Thus, for each i in the ONE treatment we cannot observe what p_{ij}^{ONE} would be if that subject had participated in the IND treatment. Instead, our null hypothesis is that each p_{ij}^{ONE} is randomly drawn with replacement from the population of 84 p_{ij}^{IND} values we observe in the IND treatment. Formally, if P_j is the set of all p_{ij}^{IND} observed in the IND treatment and $\nu(\cdot)$ is the uniform distribution over $(P_j)^{93}$, then the cdf of x_j is given by the mixture distribution $F(x_j|\nu) = \sum_{p'_j} \nu(p'_j)\Phi(x_j|\mu'_j, \sigma'_j)$, where each μ'_j and σ'_j are derived from p'_j , as above. Since $|(P_j)^{93}|$ is large, we estimate the cdf by randomly sampling 100,000 points from $(P_j)^{93}$ and using that sample to generate an estimate of $F(x_j|\nu)$.²⁶

The bottom two rows of Table VIII report $F(x_j|\nu)$ and echo the comparisons with the SIM treatment: In PM questions subjects choose the dominant significantly more often in the ONE treatment, while in the RS questions subjects jump from choosing the risky option less often to more often as it becomes more attractive, though statistical significance in the RS questions is clearly lower.

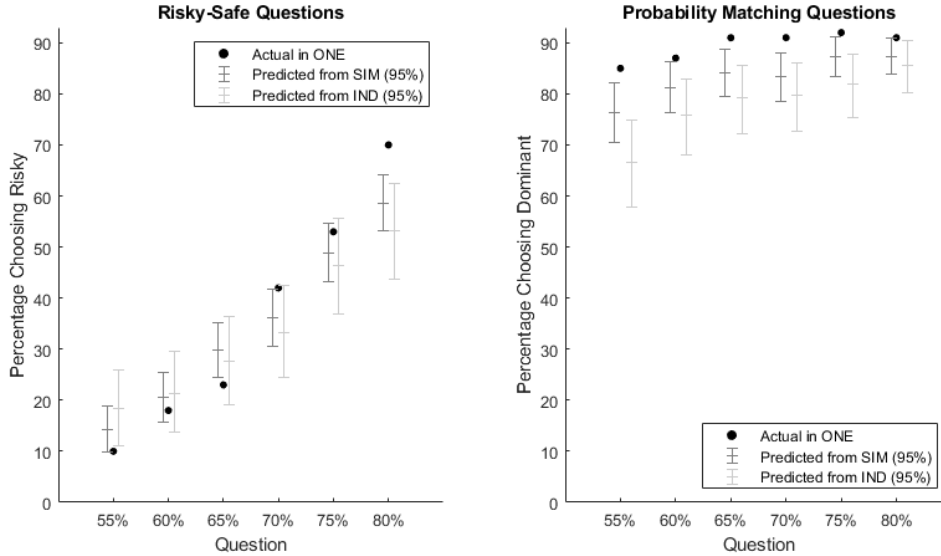


Figure VI: Number of subjects choosing the risky or dominated option in ONE, compared to the numbers predicted by mixing frequencies in the IND and SIM treatments. Vertical bars show 95% prediction ranges.

The combined results are visualized in Figure VI. Overall, we clearly reject the hypothesis that the frequency of PM choices in the ONE treatment is identical to that

²⁶Running this simulation multiple times reveals that the estimates of $F(x_j|\nu)$ differ by less than 0.001 (at most) across simulations.

of the SIM or IND treatments (right panel). Statistical results for the RS questions are mixed (left panel), but behavior in the ONE treatment is clearly more extreme—shifting from too low to too high—and not well explained by the mixing probabilities in either the SIM or IND treatments.

	PM	RS	SC	SUMP	PM55	PM80
RS	0.46***					
SC	0.77***	0.44***				
SUMP	0.40***	0.25**	0.56***			
SUDS	0.14	0.070	0.34***	0.33***		
SC55					0.71***	
SC80						0.15

Table IX: Pairwise Correlations in Individual Mixing in CORR treatment controlling for Risk-Attitudes

Notes: *** indicates significance at 1% level.

B.2. Risk Attitudes in CORR and IND treatments

Recall that subjects made two distinct risky investment decisions. We analyze these decisions in light of our randomization typology to highlight how mixers and non-mixers can differ “out-of-sample.” In both risky investment decisions, we endow subjects with \$10, any portion of which they could invest in a risky project. If the project is successful, which occurs with probability p , the amount invested is multiplied by R and paid to the subject. If the project is unsuccessful, the amount invested is lost. In either case, subjects keep the portion of the endowment they chose not to invest. The parameters used in the two risky investment tasks are $(p = 0.5, R = 2.5)$ and $(p = 0.4, R = 3)$.

	IND	CORR
Never Mix	0.882	0.967
Sometimes Mix	0.864	0.886
Always Mix	0.734	0.855

Table X: Correlation between two Risk Measures by Randomization Type

We find that the two measures of risk preferences exhibit higher correlation for our subjects who mix *less*. Individuals who always mix (in all four domains) have

the least-correlated risk measures, while those who never mix (in any of our four domains) exhibit the highest correlation.²⁷ While we view this as an exploratory exercise, we interpret this as suggestive evidence that identifying randomization types can improve other measurement exercises and out-of-sample predictions.

Table XI reports Spearman correlations between tendency to mix in various domains and a binary variable (“Risk Less”) that takes value one if a subject chose to invest less than their whole endowment in at least one of the two investment tasks (Block IV) and zero otherwise.²⁸ Table XI shows that risk averse subjects are more likely to mix in all four domains and more likely to do so for a larger number of decision problems in the PM and RS domains. This relation is strong and holds even for PM80 and RS80 questions, for which the likelihood of mixing is among the lowest among all questions considered. The relationship between risk preferences and randomization behavior gives further evidence that randomization is a stable individual trait.

	Indicator if mixing in				
	PM (PM80)	RS (RS80)	SC (SC80)	SUMP	SUDS
Risk Less	0.17*** (0.14***)	0.21*** (0.17***)	0.16** (0.01)	0.19**	0.07
	# of questions in which a subject mixes				
	PM	RS			
Risk Less	0.19***	0.19***			

Table XI: Correlation b/w Mixing Behavior and Risk Attitude in the Main Experiment

Notes: We report the pair-wise correlations between mixing behavior in different domains and an indicator taking the value of one if a subject chose to invest less than their whole endowment in at least one investment task. ***, **, and * denote significance at 1%, 5%, and 10% level, respectively.

B.3. Mixing Definition

In the main text, we say that a subject “randomizes” if they choose less than 90% of the same choices in a given decision problem. We consider them a “mixer” in a given domain if they randomize on at least one question in the domain. We consider

²⁷We exclude one subject who risked their full endowment in one question and risked nothing in the second question.

²⁸Under expected utility, risking less than the full endowment indicates that the subject is clearly risk averse.

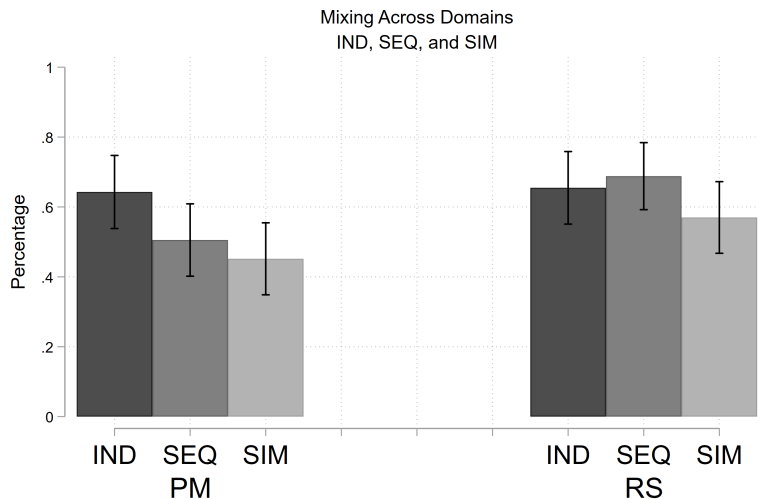


Figure VII: Mixing Behavior in IND, SEQ, and SIM Treatments for All Sequences

alternatives to this definition by tightening the definition per question, and by relaxing the definition per domain. All of the results reported are for the IND treatment.

In Figure VIII, we define randomization on a given question as choosing at least one different choice per decision problem. We define randomization per domain as mixing on a single decision problem in that domain, as in the main text. Naturally, randomization rates are slightly higher per question, but the same general trends emerge. Nevertheless, we see the same correlations across domains, as reported in Table XII.

	PM	RS	SC	SUMP
RS	0.57***			
SC	0.55***	0.39***		
SUMP	0.60***	0.45***	0.74***	
SUDS	0.44***	0.27**	0.42***	0.44***

Table XII: Pairwise Correlations in Individual Mixing in the IND Experiment: Alternative Definition 1

Notes: We report pairwise correlations between indicator variables indicating whether a subject mixed in each of our decision environments. *** indicates significance at 1% level.

In Figure IX, we define randomization on a given question as choosing at least one

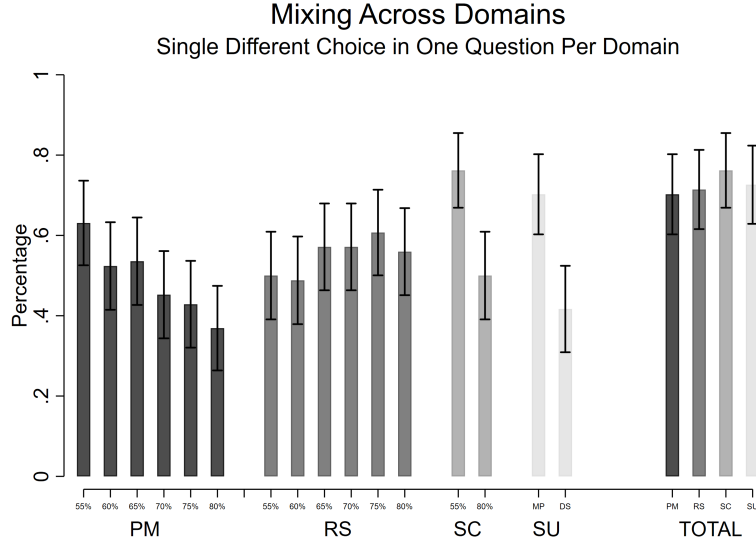


Figure VIII: Mixing Behavior Across Domains: Alternative Definition 1

different choice per decision problem, as above. We define randomization per domain as mixing on at least two decision problems in that domain. Randomization per domain mostly decreases for SC and SU, where there are only two decision problems. Nevertheless, we see the same correlations across domains, as reported in Table XIII.




Figure IX: Mixing Behavior Across Domains: Alternative Definition 2

	PM	RS	SC	SUMP
RS	0.66***			
SC	0.63***	0.50***		
SUMP	0.59***	0.55***	0.60***	
SUDS	0.44***	0.38***	0.46***	0.44***

Table XIII: Pairwise Correlations in Individual Mixing in the IND Experiment: Alternative Definition 2

Notes: We report pairwise correlations between indicator variables indicating whether a subject mixed in each of our decision environments. *** indicates significance at 1% level.

Finally, in Figure X, we define randomization on a given question as in the main text, choosing less than 90% of the same choices in a decision problem. We define randomization per domain as mixing on at least two decision problems. Similar to the above definition, the main decrease in mixing comes from the SC and SU games. Nevertheless, we see the same correlations across domains, as reported in Table XIV.

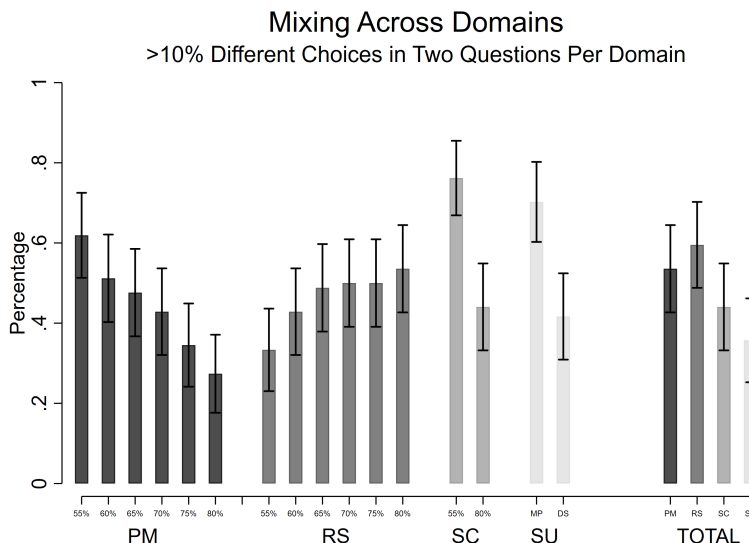


Figure X: Mixing Behavior Across Domains: Alternative Definition 3

B.4. The Intensive Margin of Mixing

In Figures XI–XIII we show histograms of the number of choices made by subjects in each question in the IND treatment.

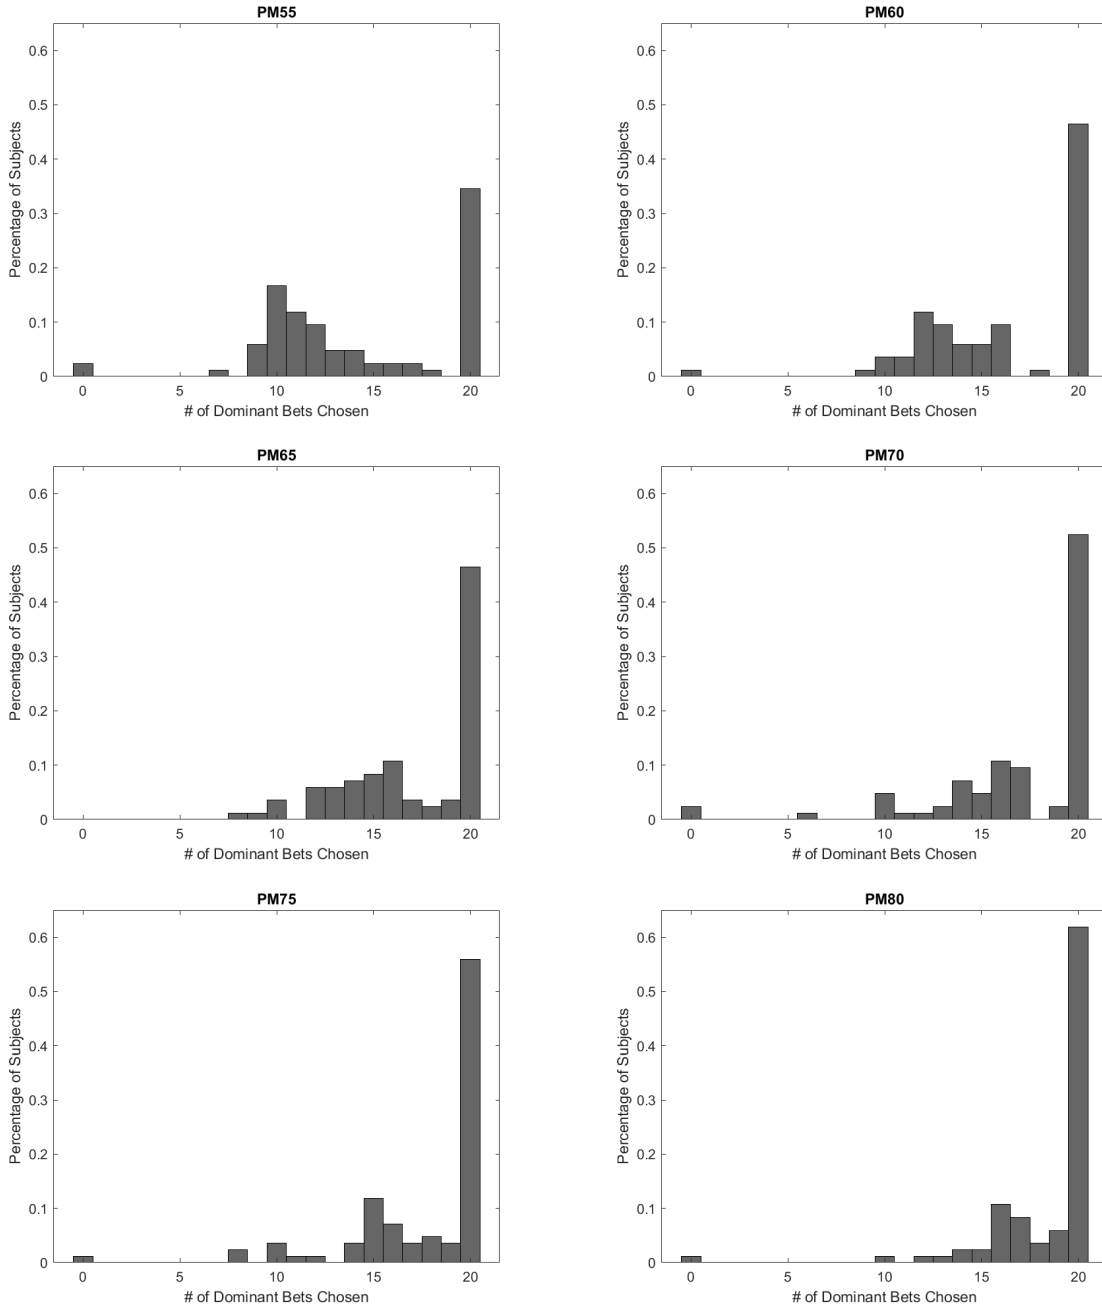


Figure XI: Histograms of the frequency of choices in the PM questions.

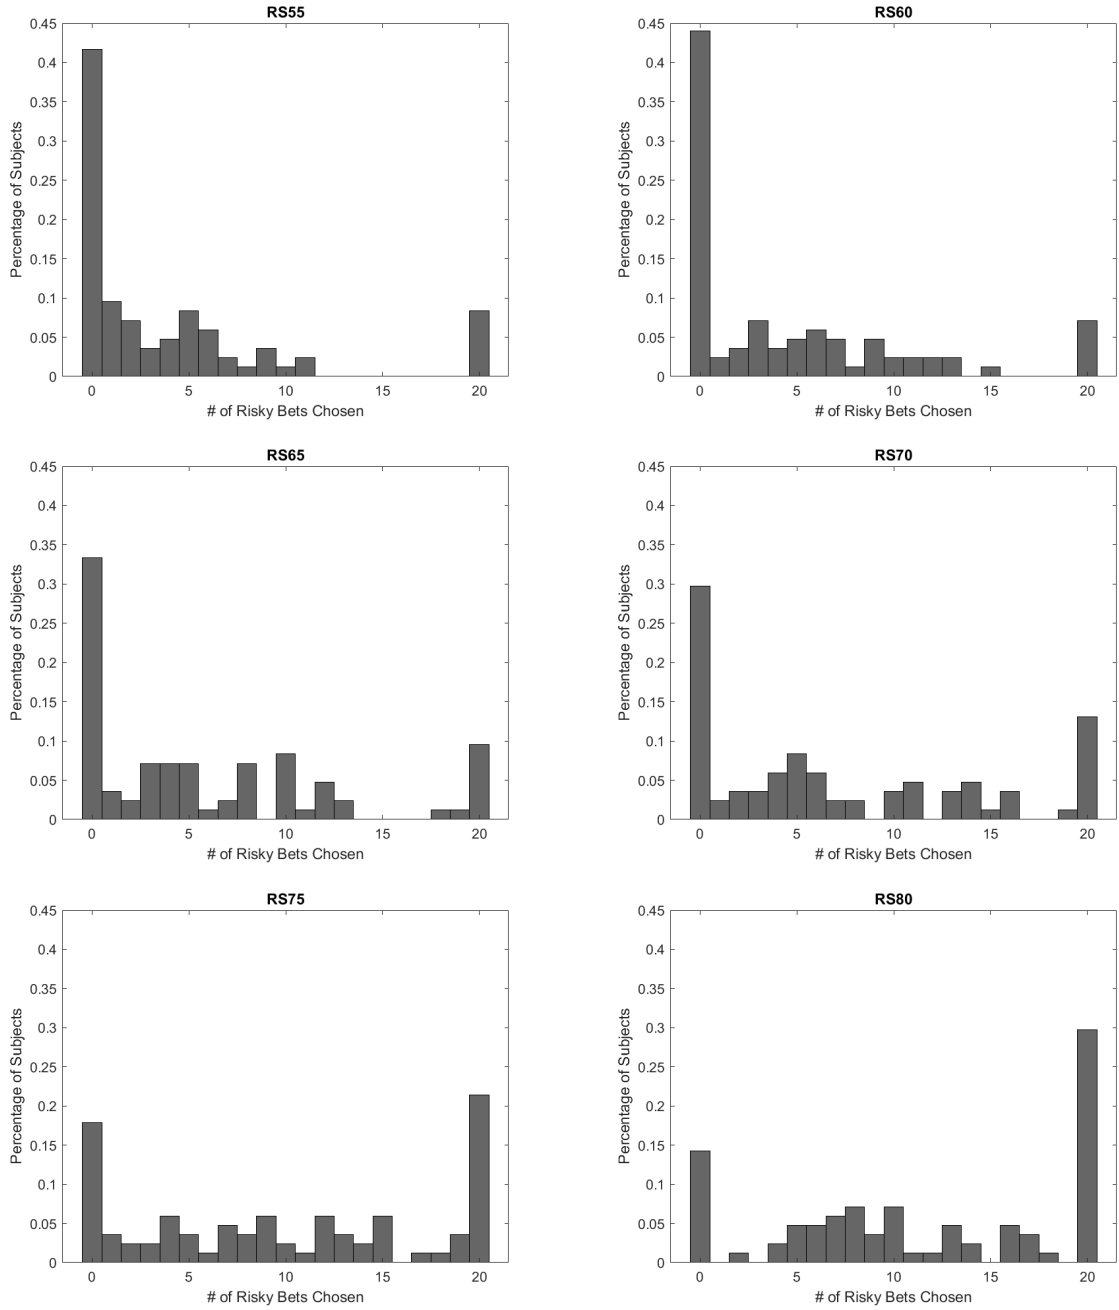


Figure XII: Histograms of the frequency of choices in the RS questions.

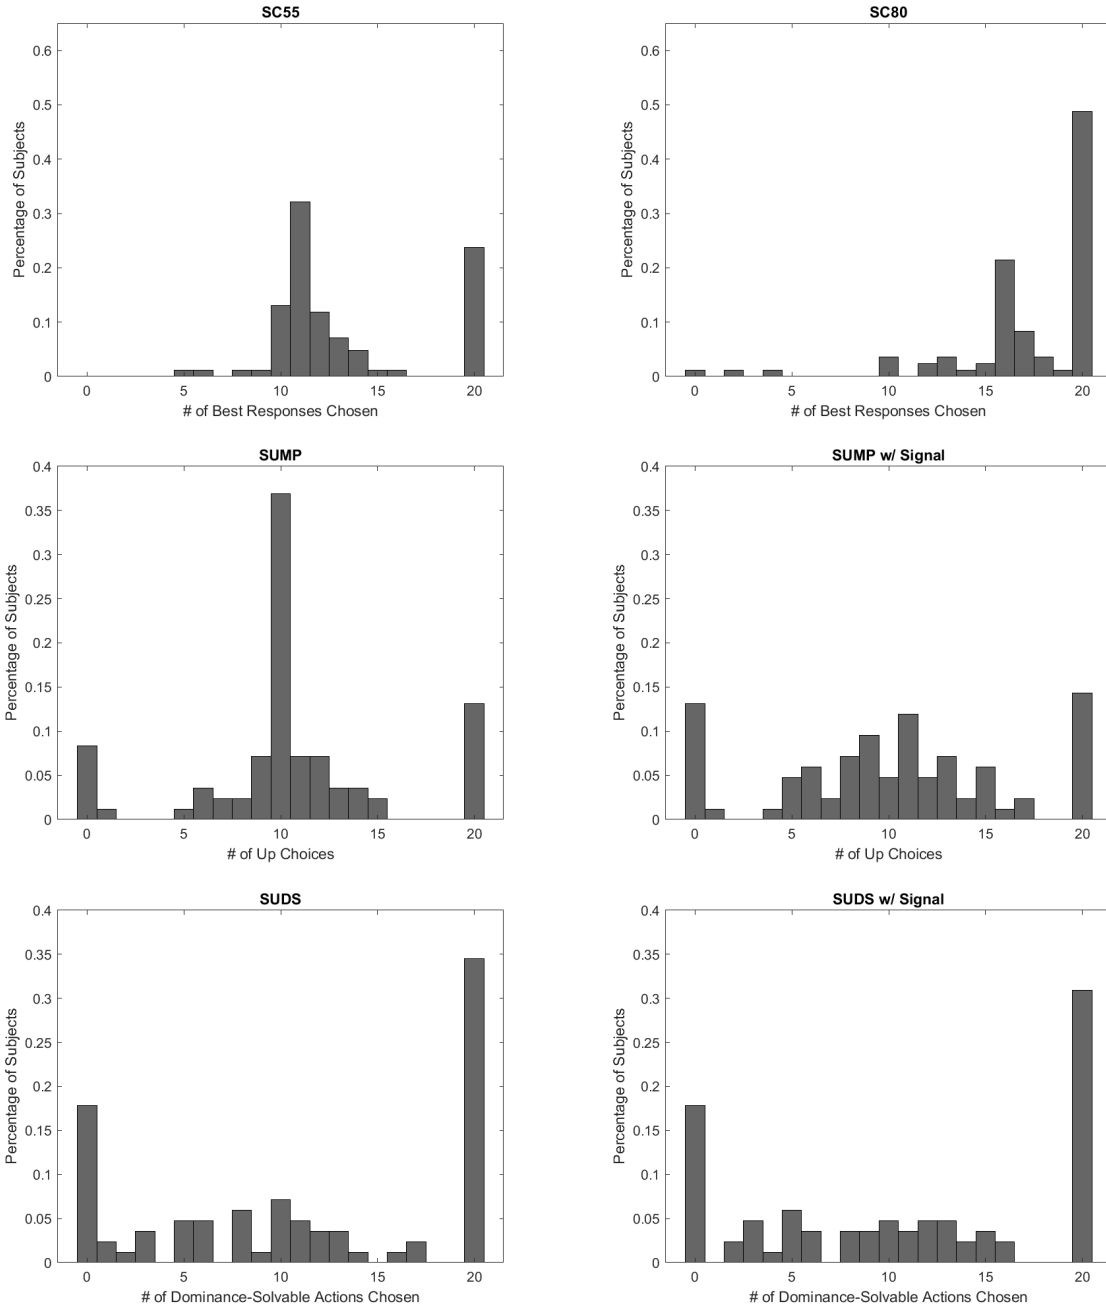


Figure XIII: Histograms of the frequency of choices in the games.

	PM	RS	SC	SUMP
RS	0.59***			
SC	0.54***	0.44***		
SUMP	0.51***	0.55***	0.54***	
SUDS	0.42***	0.43***	0.51***	0.46***

Table XIV: Pairwise Correlations in Individual Mixing in the IND Experiment: Alternative Definition 3

Notes: We report pairwise correlations between indicator variables indicating whether a subject mixed in each of our decision environments. *** indicates significance at 1% level.

C. THEORIES OF MIXING

In this appendix we describe several candidate theories that can explain mixing behavior, particularly in PM questions where mixing is stochastically dominated.

To begin, we briefly introduce our theoretical framework. Subjects face many choice problems, each of which has 20 replicates (or, one replicate in the ONE treatment). Each choice problem is given by $D_i = \{f_i, g_i\}$, where f_i and g_i are acts that map Bingo ball draws into monetary payments. Letting $B = \{1, \dots, 20\}$ be the set of possible ball draws and $X = \{\$2, \$1, \$0\}$ represent the three possible monetary payments, an act is thus a function $f_i : B \rightarrow X$. Each ball is drawn with probability $1/20$, so we can also think of f_i as a the lottery it induces over X .

In some examples we imagine a “smaller” version of our experiment with $n < 20$ replicates and n possible ball draws; the adjustment in notation for those cases should cause no confusion.

We use a shorthand notation for acts: $x_k y$ means $x \in X$ is paid if $b \leq k$, and $y \in X$ is paid otherwise. Thus, $2_{13}0$ is the act that pays \$2 if the ball drawn is 1–13, and \$0 if the ball drawn is 14–20. We simply write 1 for the act that pays \$1 regardless of b . With this notation, our PM questions are of the form $\{2_k 0, 0_k 2\}$ for $k \in \{1, \dots, 16\}$. Notice that $2_k 0$ stochastically dominates $0_k 2$ (when viewed as lotteries), but does not dominate it state-by-state. Our RS questions are of the form $\{2_k 0, 1\}$, which has no dominance relationship.

Let j index the 20 replicates of the i th decision problem, and $a_{ij} \in \{f_i, g_i\}$ be the subject’s actual choice on replicate j . The vector $a_i = (a_{i1}, \dots, a_{i20})$ represents choices made on all twenty replicates. In the IND treatment there is a ball draw b_j for each replicate j . If replicate j of problem i is chosen for payment then the subject is paid $a_{ij}(b_j) \in X$. In the CORR treatment there is only one ball draw, denoted $b_1 \in B$, and the subject is paid $a_{ij}(b_1)$ when replicate j is chosen.

Because we are interested in mixing across replicates, we assume the subject has a preference over all 20 replicates given by \succeq over the various possible vectors a_i . Their preference on a single replication is then given by \succeq_0 over acts themselves. For example, a subject who has $f_i \succeq_0 g_i$ might then have $(f_i, f_i, \dots, f_i) \succeq (g_i, f_i, \dots, f_i)$. Or, if the subject prefers to mix, then perhaps some mixed vector would be preferred over (f_i, f_i, \dots, f_i) .

We say that \succeq respects *replicate dominance* if $a_i \succeq a'_i$ whenever, for every j ,

$a_{ij} \succeq_0 a'_{ij}$. An implication of replicate dominance is that if $f_i \succeq_0 g_i$ on the individual replicates then (f_i, f_i, \dots, f_i) must be preferred over every other vector a'_i .²⁹ Thus, replicate dominance rules out mixing on any problem.

Similarly, \succeq respects *stochastic dominance* if $a_i \succeq a'_i$ whenever, for every j , a_{ij} stochastically dominates a'_{ij} . Thus, respecting stochastic dominance means that the subject can never mix in any PM question, but can possibly mix in RS questions. Notice that if \succeq_0 always selects stochastically dominant options then respecting replicate dominance implies the subject also respects stochastic dominance.

In our experiment we observe mixing on PM questions, even though we do see strong evidence in the ONE treatment that \succeq_0 selects the stochastically-dominant option for the vast majority of subjects. Thus, for a large number of subjects, \succeq respects neither replicate dominance nor stochastic dominance. We therefore seek a theory in which (1) replicate and stochastic dominance are not respected, (2) \succeq_0 selects stochastically dominant options, and (3) these patterns hold true in both the IND and CORR frameworks. We now review a handful of models and show that none satisfactorily satisfy all three requirements.

C.1. Preferences Over Reduced Lotteries

As discussed in the paper, mixing implies convex preferences over reduced lotteries, and mixing in PM questions requires violations of dominance. Here we describe in further detail two theories in this domain.

C.1.a. Regret Aversion

Loomes and Sugden (1982) describe a theory of regret aversion wherein a subject experiences regret in some state if an alternative choice would have yielded a higher utility index. For example, consider the PM75 question $D_i = \{f, g\}$, where $f = 2_{15}0$ and $g = 0_{15}2$, in the CORR treatment. If $L = \{b : b \leq 15\}$ obtains then choosing f on a given replicate gives *ex-post* payoff $u_2 := u(2) + R(u(2) - u(0))$, while choosing g gives $u_0 := u(0) + R(u(0) - u(2))$. The opposite payoffs occur in event $H = \{b : b > 15\}$. Loomes and Sugden (1982) assume $R(0) = 0$ and $R(\cdot)$ is non-decreasing, which implies $u_2 > u_0$. Consequently, the “regret-adjusted” payoffs of choosing f stochastically dominate

²⁹If we view f_i and g_i as lotteries, this implication is called *compound betweenness*, which is a weakening of the compound independence axiom; see Camerer and Ho (1994).

those of g , and so a subject maximizing regret-adjusted expected utility should not mix.

C.1.b. Probability Weighting

Under this theory a subject evaluates the lottery $p = (p_1, x_1; \dots; p_n, x_n)$ according to the functional

$$U(p) = \sum_i w(p_i)u(x_i),$$

where $w : [0, 1] \rightarrow [0, 1]$ is a probability weighting function that is onto and strictly increasing, and u is a strictly increasing utility index. Behaviorally, this model assumes that subjects first transform the vector of probabilities into a vector of weights (which may not sum to one) and then satisfy the independence axiom using these weighted probabilities.

Consider a PM question where the dominant bet pays off with probability $p > 1/2$. Denote the proportion of dominant bets the subject chooses by $q \in [0, 1]$. Then their overall utility is given by

$$w(qp + (1 - q)(1 - p))u(2) + w(q(1 - p) + (1 - q)p)u(0).$$

If this function is *decreasing* in q at $q = 1$ then the subject will not choose the dominant bet in every replication.³⁰ It is decreasing at $q = 1$ if

$$\frac{u(2)}{u(0)}w'(p) < w'(1 - p). \quad (1)$$

Since $u(2) > u(0)$, we get mixing only when w is highly asymmetric: very steep at low probabilities $(1 - p)$ and flat for high probabilities (p) .³¹ In that case the subject is happy to sacrifice their 20th dominant bet (which costs them $w'(p)u(2)$ on the margin) for their first dominated bet (which gains them $w'(1 - p)u(0)$ on the margin).³²

In theory it is possible to find such a function, but standard weighting functions in the literature do not feature this sort of asymmetry. For the standard Prelec weighting function ($w_P(p) = e^{-\beta(-\ln(p))^\alpha}$) the slopes are not sufficiently asymmetric

³⁰For simplicity we assume here that q can take any value in $[0, 1]$.

³¹Decreasing at $q = 1$ is sufficient to show mixing if the objective is concave. Roughly speaking, for inverse-S weighting functions the objective is concave for all $q \in [0, 1]$ as long as $w(\cdot)$ is not “too convex” at p .

³²We omit a common factor of $2p - 1$ on both margins.




Figure XIV: The usual Prelec weighting function (dashed) is not sufficiently asymmetric to generate mixing. The modified version (solid) is, but its inflection point is too high to match empirical estimates.

(see the dashed curve in Figure XIV). But the weighting function $w(p) = 1 - w_P(1 - p)$ (which simply takes the Prelec function and flips both the x - and y -axes; see the solid curve in Figure XIV) is asymmetric enough to generate an interior maximum. For example, with $\alpha = 0.6$, $\beta = 1.6$, and $u(x) = x^{0.5}$, the optimal mix in the $p = 0.80$ question is $q = 17/20$. Similar calculations show that this function can also predict mixing in the RS questions.

Unfortunately, any inverse-S-shaped function that can generate mixing (steep for low p and flat for high p) is necessarily going to have an inflection point that's too high to match empirical estimates. Wu and Gonzalez (1996) estimate that the inflection point should be at or below 0.40—meaning $w'(\cdot)$ is increasing for $p > 0.40$ —but mixing in PM55 implies that $w'(0.45) > w'(0.55)$.

Source preference theories (Tversky and Fox, 1995; Abdellaoui et al., 2011, e.g.) are a generalization of probability weighting that allow for different weighting functions on different sources of uncertainty. For example, which replicate is paid and which ball is drawn. In applications, however, source functions differ only if events have unknown probability, as in the Ellsberg paradox. In our experiment all uncertainty is objective, so we apply only a single probability weighting function.

C.2. Preferences Over Two-Stage Lotteries

C.2.a. Perturbed Utility Models

Applying the perturbed utility model of Allen and Rehbeck (2019) to our setting, a subject facing decision problem $\{f_i, g_i\}$ who picks f_i in k out of 20 replicates receives utility

$$\frac{k}{20}U(f_i) + \frac{20-k}{20}U(g_i) + V\left(\frac{k}{20}\right),$$

where $U(\cdot)$ is any utility for the underlying lotteries—it need not be consistent with any particular model like expected utility—and $V(k)$ has a unique maximizer. For exposition, assume k is continuous and V is differentiable so that the optimal k^* is given by the first-order condition

$$-V'\left(\frac{k}{20}\right) = U(f_i) - U(g_i).$$

In other words, the subject balances their marginal preference for randomization with the utility difference between the two options.

In the SEQ treatment a subject making their choice on the j th replicate chooses a plan to select f_i in $k_j \in \{0, \dots, n-j\}$ remaining replicates, giving a utility of

$$\frac{k}{n-j}U(f_i) + \frac{n-j-k}{n-j}U(g_i) + V\left(\frac{k}{n-j}\right).$$

Under the differentiability assumption, this is maximized at the k^* for which

$$-V'\left(\frac{k}{n-j}\right) = U(f_i) - U(g_i).$$

Thus, this model does not predict a drop in mixing frequency in the SEQ treatment. It is possible, however, that a different formulation of the $V(\cdot)$ function could predict a difference; a formulation in which not only the proportion matters but also the number of choices.

C.2.b. Siegel's Perturbed Utility Model

Siegel (1961, Model II) proposes a model of decision-making designed expressly to predict probability matching. According to this theory the subject experiences greater

utility for predicting the less-likely event, and also receives positive utility for varying their choice across replications. The latter is described as a direct preference for avoiding monotonous repetition, and is not related to the variance in outcomes (such as with risk-seeking expected utility).

To illustrate, consider the PM75 question, where betting f has a 15/20 chance of paying off. Normalize to 1 the marginal utility of a correct prediction of the more-likely event, let $\alpha \geq 1$ be the marginal utility of a correct prediction of the less-likely event, and let $\beta \geq 0$ be the marginal utility of variance. The overall expected utility of choosing f k times out of 20 (in either IND or CORR) is then given by

$$U(f_k g) = \frac{15}{20} \frac{k}{20} + \alpha \frac{5}{20} \frac{20-k}{20} + \beta \frac{k}{20} \frac{20-k}{20}.$$

Assuming $c > 0$ and ignoring integer constraints, this is maximized at

$$k^* = 10 + \frac{15 - 5\alpha}{2\beta}.$$

Without knowledge of α and β , this model can predict any mixing behavior in both the IND and CORR treatment. The exact probability matching result ($k^* = 15$) obtains whenever $15 - 5\alpha = 10\beta$. In fact, an increased utility for predicting rare events ($b > 1$) is not needed for this prediction, as $\alpha = \beta = 1$ is one parameterization that leads to $k^* = 15$.

Our purpose in comparing various models is to identify possible underlying causes for mixing behavior. To that end we find this model somewhat unenlightening because it essentially assumes the result: subjects mix because they have a direct preference for varying their choices. If we shut down this direct preference for variation in choices (by setting $\beta = 0$) then the model predicts that the subject will choose their most-preferred choice in all 20 replicates (that is, f if $\alpha < 3$ and g if $\alpha > 3$).

C.2.c. u-v Preferences and Utility for Gambling

Several authors (Neilson, 1992; Schmidt, 1998; Diecidue et al., 2004) have described versions of a model where preferences satisfy expected utility on the interior of the simplex with utility index u , but certain payments are evaluated by a different function $v \neq u$. This can be used to explain a disproportionate preference for certainty ($v > u$) or an explicit preference for gambling ($u > v$).

In our CORR treatment we view the state space to be $C \times B$ —capturing both the choice of which replicate is paid and which ball is drawn. With that view, no PM question offers certainty, as both bets have some chance of not paying off. Since these models assume expected utility away from certainty, they cannot explain mixing in the PM questions.

Suppose instead we view the problem to be a two-stage lottery, where C is chosen first and then B is chosen second. Choosing f on all 20 replicates does guarantee certainty *in the first stage*, but not the second. We could apply u - v preferences to the first stage alone, in which case mixing can be predicted. Indeed, mixing is then equivalent to a violation of replicate dominance (defined in the introduction of this appendix).

C.3. Mistakes, Biases, and Heuristics

Here we review in more detail those models discussed in the text.

C.3.a. Gambler’s Fallacy: Expected Utility with Negative Correlation

The gambler’s fallacy is the mistaken belief in negative serial correlation. Consider a probability matching question $\{2_k 0, 0_k 2\}$ where $k > 10$. Roughly, a subject who chooses $2_k 0$ (or, bets $b \leq k$) several times in a row might feel that $b > k$ is “due” on the next replicate. Now, ball draws are not observed sequentially, but this belief in negative correlation can still drive subjects to exhibit mixing in PM questions.

To show this formally, fix $k > 10$ and let $L = \{b : b \leq k\}$ be the event that bet $2_k 0$ pays off, and $H = \{b : b > k\}$ be the event that bet $0_k 2$ pays off. The objective probability of event L is $p = k/20$. A subject exhibiting the gambler’s fallacy wrongly believes events L and H are negatively correlated. Let $p(L|L)$ denote the subject’s probability of L on some replicate $j > 1$ given that L occurred on $j - 1$. A fully rational subject would have $p(L|L) = k/20$, but instead we model the subject’s belief as

$$p(L|L) = \alpha 0 + (1 - \alpha)p,$$

where $\alpha \in [0, 1]$ is a simple way of capturing the subject’s degree of gambler’s fallacy. We refer to this belief as α -negative correlation, where $\alpha = 0$ represents the objectively correct belief.

State	Probability (α -Negative Correlation)		ff Pays	fg Pays
LL	$p[\alpha 0 + (1 - \alpha)p]$	$= (1 - \alpha)p^2$	1	$\frac{1}{2}$
LH	$p[\alpha 1 + (1 - \alpha)(1 - p)]$	$= p - (1 - \alpha)p^2$	$\frac{1}{2}$	1
HL	$(1 - p)[\alpha 1 + (1 - \alpha)p]$	$= p(1 - p) + \alpha(1 - p)^2$	$\frac{1}{2}$	0
HH	$(1 - p)[\alpha 0 + (1 - \alpha)(1 - p)]$	$= (1 - \alpha)(1 - p)^2$	0	$\frac{1}{2}$

Table XV: A two-choice PM example with α -negative correlation.

To illustrate the model, suppose the subject chooses between $f = 2_k 0$ and $g = 0_k 2$ only twice (instead of 20 times). The four relevant states of the world and their corresponding probabilities are then given in the first two columns of Table XV. A subject considers two possible betting strategies: betting f both times (ff) and betting f then g (fg). Assuming expected utility and normalizing $u(2) = 1$ and $u(0) = 0$, the expected payoff of each strategy for each state—taking into account that each bet has a $1/2$ chance of being paid—is shown in the right two columns of Table XV.

From the table we calculate the following expected utilities for each strategy:

$$Eu(ff) = p - \alpha \frac{1}{2} (p^2 - (1 - p)^2), \text{ and}$$

$$Eu(fg) = p - (1 - \alpha) \frac{1}{2} (p^2 - (1 - p)^2).$$

From these we can see that mixing (fg) is preferred if and only if $(1 - \alpha) < \alpha$, or $\alpha > 1/2$. Our data show that mixing propensity changes with p , but, at least for this simple specification, that is not predicted.

If there are $n > 2$ replicates then the strategy of alternating bets ($fgfgfg\dots$) will continue to be optimal for large enough α .³³ Here we assume the subject views draws as a Markov process, with each draw affected only by the previous draw. If we assumed a more complex correlation structure we could predict more complex patterns of mixing, such as $fffgfffg$.

For RS questions, a subject prefers f (which pays \$2 with probability p) over g (\$1 for sure) in a single decision if and only if $u(1) < p$, where $u(1) \in (0, 1)$ represents

³³To see this, consider $\alpha = 1$. The subject believes the sequence $LHLHLH\dots$ will occur with probability $p > 1/2$, and $HLHLHL\dots$ will occur with probability $1 - p < 1/2$. Thus, the alternating bet strategy is strictly optimal. By continuity this must be true for all α in some neighborhood of one. In fact, numerical calculations suggest that the threshold α^* may even decrease in n , though we have yet to prove this claim.

the subject's risk aversion. With two decisions, ff pays as above, while fg gives an expected utility of $(1/2)p + (1/2)u(1)$. Thus, fg is preferred over ff whenever $u(1) > p - \alpha[p^2 + (1-p)^2]$. In other words, mixing can occur whenever $p > u(1) > p - \alpha[p^2 + (1-p)^2]$. The larger is the value of α , the wider the range of risk preferences in which we predict mixing. Mixing is never predicted, however, if $u(1) > p$, because such a subject always prefers gg (with sure payoff of $u(1)$) over fg (with expected utility $(1/2)p + (1/2)u(1)$). Thus, we should expect to see some mixing in the questions

Mixing in the SEQ treatment under this model is identical to the IND treatment because information about which replicate is paid would not affect the belief in correlation across ball draws. The lack of mixing in the ONE treatment is also predicted, since only a single draw is realized.

In the CORR treatment, however, mixing is not predicted. There is only one ball drawn and so there are only two possible states of the world: L and H . Negative correlation cannot affect beliefs, and so this model reduces to the standard expected utility framework in which mixing is strictly dominated in PM questions and depends on whether $u(1) > p$ in RS questions. Thus, the gambler's fallacy cannot explain mixing in the CORR treatment.

Rabin (2002) models the gambler's fallacy as a subject who believes draws are made without replacement. This particular form of negative correlation does *not* predict mixing in our IND treatment. For example, in the PM75 question (where $s = 15$) the subject would believe that that f will pay off in exactly 15 bets and g will pay off in exactly 5, but all orderings of those outcomes are equally likely. Thus, from an *ex-ante* perspective, there is no reason to believe that any one pattern of outcomes is more likely than another, and so there is no reason to generate any particular pattern of bets. Our version of negative correlation, however, does produce expected patterns and, thus, optimal betting patterns in the IND treatment. Rabin and Vayanos (2010) propose a model of the gambler's fallacy much closer to ours which does predict mixing in the PM domain for IND, but not CORR.

C.3.b. Modal Count Heuristic

According to this heuristic the decision-maker focuses on the *number* of times an event will occur, but not the order of events, and their objective is to maximize the number of "correct" bets made. For example, in the PM60 question, the subject wrongly focuses

Event	Probability	Num. of L	Total Prob.
LLL	$8/27$	3	$8/27$
LLH	$4/27$		
LHL	$4/27$	2	$12/27$
HLL	$4/27$		
LHH	$2/27$		
HLH	$2/27$	1	$6/27$
HHL	$2/27$		
HHH	$1/27$	0	$1/27$

Table XVI: Modal Count Example: Possible Events and Corresponding Probabilities

on the fact that the dominant bet is most likely to pay off in 12 of the 20 replicates, so they choose the dominant bet 12 times. They do so because they wrongly believe this maximizes the chance of *all* bets paying off.

To illustrate, consider the simpler case of three replicates and three Bingo balls ($B = \{b_1, b_2, b_3\}$). Let $f = 2_20$ be a bet on $L = \{b_1, b_2\}$, and $g = 0_22$ be a bet on $H = \{b_3\}$. Since $n = 3$ there are $2^3 = 8$ possible payoff-relevant events, which we enumerate in Table XVI. In this case, two L s and one H is the most likely outcome count, with a total probability of $12/27$, so the subject bets f twice and g once. What they fail to realize is that the order of their bets matters and in fact their true probability of getting all three bets correct is only $4/27$, not $12/27$.

In the CORR treatment, however, there are only two possible outcomes: L obtains for all twenty replicates, or H obtains for all twenty replicates. Thus, a subject focused on the modal number of outcomes should never mix.

C.3.c. Regret & Convex Costs of Mistakes

According to this theory, the subject has a convex cost of “mistakes,” where a mistake is simply a bet that doesn’t pay off. Choosing $f = 2_k0$ all twenty times in the CORR treatment opens the possibility that all twenty bets turned out to be wrong, *ex post*. Mixing reduces the maximum number of mistakes the subject might make.

Formally, the subject’s preference over mixtures is represented by the menu-

dependent utility function

$$U(f_k g | \{f, g\}) = \frac{k}{n}v(f) + \frac{n-k}{n}v(g) - \alpha \frac{1}{20} \sum_{b \in B} \left(w\left(\frac{k}{n}\right) \max\{g(b) - f(b), 0\} + w\left(\frac{n-k}{n}\right) \max\{f(b) - g(b), 0\} \right),$$

where $B = \{1, \dots, 20\}$, $v(\cdot)$ represents preferences over degenerate acts in X^B , $\alpha \geq 0$ is an individual-specific scale parameter, and $w(\cdot)$ is an increasing and weakly convex function satisfying $w(0) = 0$. The summation term counts for each state the fraction of times the subject made the “wrong” choice in that state, which is then weighted by the convex function $w(\cdot)$ and multiplied by the payoff magnitude of the mistake. Convexity captures the idea that the decision maker finds it especially undesirable to have states in which they have made many incorrect choices. Thus, they would gladly add mistakes in states where they have relatively few in order to reduce mistakes in states where they have many.

To see how this preference generates probability matching, consider the probability matching decision problem $D_i = \{f, g\}$, where $f = 2_k 0$ and $g = 0_k 2$, with $k > 10$. If the decision maker picks f in s replications (so, picks $f_s g$) and $w(x) = x^2$ then the cost term becomes

$$-\alpha \frac{1}{20} \left(k \left(\frac{n-s}{n} \right)^2 (2-0) + (20-k) \left(\frac{s}{n} \right)^2 (2-0) \right).$$

In words, there are k states in which f is the better choice but g is chosen in $\frac{n-s}{n}$ of the replications, and there are $20-k$ states in which g is the better choice but f is chosen in $\frac{s}{n}$ of the replications. Maximizing over s gives the solution

$$s^* = \frac{k}{20}n,$$

which is exactly the probability matching prediction. This decision maker faces a tension between choosing the act with the higher base value (by comparing $v(f)$ to $v(g)$) and performing probability matching to reduce the cost of mistakes. Individuals with a higher value of α will lean more toward probability matching, while individuals with $\alpha = 0$ will choose the more-preferred act in every replication.

In the IND treatment, however, we can show that the distribution of the number of mistakes shifts up (in the sense of first-order stochastic dominance) whenever a

bet of f is replaced by a bet of g . Intuitively, betting on g increases the chance of a mistake on this replicate, and, regardless of the draws of the balls, is unrelated to the number of mistakes made in other replicates. Thus, the expected cost of mistakes necessarily increases for any w . Betting f all 20 times maximizes the expected payoff and minimizes the cost of mistakes, making it the predicted choice regardless of w and α .

C.3.d. Responsibility Aversion

Dwenger et al. (2018) also find evidence of mixing in repeated binary choices without a dominant option. They argue that this may be due to subjects delegating their decisions to a randomization device so that they can avoid being directly responsible for their outcome, as in Bartling and Fischbacher (2012). In an earlier working paper Dwenger et al. propose a formal theory of responsibility aversion for lottery choices that is tailored to their setting. That theory doesn't port well to our experiment, but it is possible that an alternative theory of responsibility aversion could explain mixing in our domains. What such a theory would look like remains an interesting open question.

C.3.e. Irrational Diversification

In this theory the subject maximizes expected utility but incorrectly believes they will be paid for all choices, rather than one randomly-selected choice. We will show that this can lead to rational mixing in the correlated treatment, but not in the independent treatment.

The intuition is as follows. Suppose $D_i = \{f, g\}$, where $f = 2_{15}0$ and $g = 0_{15}2$, so f is the dominant choice. To illustrate, let $n = 2$. Choosing $a_i = (f, f)$ in the CORR treatment gives the subject a 3/4 chance of \$4 and a 1/4 chance of \$0. But $a'_i = (f, g)$ gives the subject \$2 for sure. The bet g offers a perfect hedge in case f does not pay off. A sufficiently risk averse subject will therefore choose a'_i .

In the independent treatment, $a_i = (f, f)$ gives a 9/16 chance of \$4 and a 1/16 chance of \$0, while $a'_i = (f, g)$ gives a 3/16 chance of \$4 and a 3/16 chance of \$0. The remaining probability in both is on \$2. In this case g does not offer a hedge against losing in f since the two bets pay off independently. Here, a'_i is stochastically dominated by a_i and should never be chosen.

C.3.f. Obvious Dominance

As there are several levels of randomization, there are several possible notions of dominance. The three relevant notions are:

- a_i C -dominates a'_i if, for all j , $a_{ij} \geq_0 a'_{ij}$.
- a_i C -stochastically dominates a'_i if, for all j , a_{ij} stochastically dominates a'_{ij} .
- a_i $C \times B^n$ -dominates a'_i if, for all j and b_j , $a_{ij}(b_j) \geq a'_{ij}(b_j)$.

We see a substantial amount of mixing in the IND and CORR treatment, indicating that C -stochastic dominance is violated. But mixing is reduced in the sequential treatment. We now explore conditions that guarantee no mixing in the SEQ treatment.

Li (2017) shows how, in allocation settings, moving from a simultaneous-move auction (such as the sealed-bid second-price auction) to a sequential-move auction (such as the English clock auction) can make the dominance property of truth-telling more “obvious” to the bidder. The informal intuition is that in the clock auction the bidder only needs to consider the current clock price—should she stay in or out—whereas in the sealed-bid auction she should consider all possible highest bids of her opponent. Li (2017) formalizes this by strengthening dominance to a comparison of worst-case payoffs of the dominant plan to best-case payoffs of the considered deviation; if the worst-case payoff of the dominant plan is preferred to the best-case payoff of the deviation, then no state-by-state contingent reasoning is needed to determine dominance.

In our experiment one might expect that truth-telling (always picking the more-preferred option) is similarly more “obvious” in the SEQ treatment. We show, however, that Li’s definition of obvious dominance does not predict any treatment difference between IND and SEQ. The reason is that truth-telling is already obviously dominant in the IND treatment: The worst-case outcome under truth-telling still gives the subject their most-preferred option, so no deviation can possibly provide a better outcome. In the following we formalize this insight.

In the SEQ treatment each decision problem i can be viewed as a one-player game. Nature moves first, choosing $c_i \in C$. Then the subject has n information sets. At each information set j the subject knows $c_i \geq j$ and chooses between f_i and g_i . If $c_i = j$ then the game ends and the chosen act a_{ij} is paid (assuming $r = i$, which is revealed

at the end of the experiment). If $c_i > j$ then the subject continues to information set $j + 1$. The subject's strategy in the game is an entire plan $a_i = (a_{i1}, \dots, a_{in})$.

To give Li's definition of obvious dominance, we need to identify the first information set at which two plans a_i and a'_i differ. Formally, let $\hat{j}(a_i, a'_i) = \min\{j : a_{ij} \neq a'_{ij}\}$.

Definition 1. We can refine our original notions of C - and $C \times B^n$ -dominance to apply at each information set j :

1. a_i C -dominates a'_i at j if, for all $j' \geq j$, $a_{ij'} \geq_0 a'_{ij'}$.
2. a_i $C \times B^n$ -dominates a'_i at j if, for all $j' \geq j$ and $b_{j'} \in B$, $a_{ij'}(b_{j'}) \geq a'_{ij'}(b_{j'})$.

For each of those there is an equivalent notion of obvious dominance:

1. a_i C -obviously dominates a'_i if, for all $j \geq \hat{j}(a_i, a'_i)$ and all $j' \geq \hat{j}(a_i, a'_i)$, $a_{ij} \geq_0 a'_{ij'}$.
2. a_i $C \times B^n$ -obviously dominates a'_i if, for all $j \geq \hat{j}(a_i, a'_i)$, all $b_j \in B$, all $j' \geq \hat{j}(a_i, a'_i)$, and all $b_{j'} \in B$, $a_{ij}(b_j) \geq a'_{ij'}(b_{j'})$.

In words, the notions of obvious dominance (1) look only at the present and future information sets, and (2) compare the worst-case scenario under a_i to the best-case scenario under a'_i . In C -obvious dominance the best and worst cases are with respect to only which c_i is drawn. In $C \times B^n$ -obvious dominance the best and worst cases are with respect to both the draw of c_i and b_j .

To illustrate, suppose a subject in the SEQ treatment with $n = 5$ faces a PM question ($f_i = 2_k 0$ and $g_i = 0_k 2$ with $k \notin \{0, n\}$) and has $f_i >_0 g_i$. Consider $a_i = (f_i, f_i, f_i, f_i, f_i)$ versus $a'_i = (f_i, f_i, g_i, f_i, g_i)$. The first replicate at which these differ is at $\hat{j}(a_i, a'_i) = 3$. Under C -obvious dominance, the worst-case $j \geq 3$ under a_i is that the subject receives f_i (indeed, it is the only possible outcome). The best-case outcome under a'_i is that $j = 4$, which gives $a'_{i4} = f_i$. This is no better than the worst-case outcome under a_i , so a_i C -obviously dominates a'_i . A similar argument applies for any a'_i , so a_i is C -obviously dominant. Indeed, truth-telling ($a_i = (f_i, \dots, f_i)$) will be C -obviously dominant for any n .

For $C \times B^n$ -obvious dominance, however, $a_i = (f_i, f_i, f_i, f_i, f_i)$ does not obviously dominate a'_i . This is because for any $j \geq 3$ there some $b_j > k$ for which $f_i(b_j) = \$0$, while for any $j' \geq 3$ there is some $b_{j'}$ for which $a_{ij'}(b_{j'}) = \$2$. In other words, since f_i does not B -dominate g_i , we cannot have $C \times B^n$ -obvious dominance of a_i . An identical

argument holds for RS questions, since again the minimum payment of one choice is always strictly less than the maximum payment of the opposite choice.

But notice that the following two paragraphs hold equally true for the IND and SIM treatments. In those settings there is only one information set. For any a_i and a'_i we simply set $\hat{j}(a_i, a'_i) = 1$; otherwise the definitions above apply. And the argument that truth-telling is C -obviously dominant—but not $C \times B^n$ -obviously dominant—remains true. Thus, neither form of obvious dominance can predict mixing in the IND treatment but no mixing in the SEQ treatment.

D. NOTIONS OF DOMINANCE, MIXING, AND INCENTIVE COMPATIBILITY

For the interested reader, we provide supplementary information about various notions of dominance in our experiment, and the related notions of monotonicity that require preferences respect dominance. We explore which notions of monotonicity are violated by mixing behavior. We also discuss under which monotonicity assumptions our experiment is incentive compatible, implications for models of random preferences, and provide a modification of monotonicity—called myopic preference—that can capture mixing in the SEQ treatment.

D.1. Setup and Experimental Design

Choice objects are acts $f : B \rightarrow X$, where $B = \{1, \dots, n\}$ is the set of possible draws from a Bingo cage containing n numbered balls and $X = \{\$2, \$1, \$0\}$ is the set of possible monetary prizes.³⁴ Each ball in B is drawn with objective probability $1/n$, but we generally model choice objects as acts. We can describe f as an n -vector—such as $f = (2, 0, \dots, 1, 1)$ —to indicate the prize awarded in each state. For any two prizes $x, y \in X$ and any $k \in \{0, 1, \dots, n\}$ let $x_k y$ be the act that pays x in the first k states and y otherwise. For example, $2_{10}0$ is the bet that pays $\$2$ in states 1–10 and $\$0$ otherwise. The constant act that pays x in every state is denoted simply as x .

The subject is given m different decision problems, each of which is a choice between two acts. Denote the i th problem by $D_i = \{f_i, g_i\}$. The subject makes each of these choices n times. The subject's choice on the j th replicate of the i th problem is given by $a_{ij} \in D_i$. Let $a = (a_{ij})_{i,j} \in \times_{i=1}^m D_i^n$ be the entire matrix of choices and $a_i = (a_{i1}, \dots, a_{in}) \in D_i^n$ be the vector of choices made across the n replicates of the i th problem.

³⁴In the actual experiment $X = \{\$25, \$15, \$5\}$; we use $\{\$2, \$1, \$0\}$ only for notational convenience.

In our baseline condition one ball is drawn (with replacement) for each of the n replicates. Let $b = (b_1, \dots, b_n) \in B^n$ be the vector of all n draws. Act a_{ij} is paid based on draw $b_j \in B$. The final payment is therefore $a_{ij}(b_j) \in X$.

We employ the RPS mechanism, meaning one of the mn choices is chosen randomly for payment. The decision problem chosen (the “row” of the matrix a) is determined by a randomization device with realizations $r \in R = \{1, \dots, m\}$, and the replicate (“column”) is determined by a separate randomization device with realizations $c \in C = \{1, \dots, n\}$. Thus, the combined state (r, c) determines which problem and which replicate is paid. The announcement of $a = (a_{ij})_{i,j}$ generates an act which pays act a_{ij} in state $(r, c) = (i, j)$. And the act a_{ij} pays $a_{ij}(b_j)$ in each state $b_j \in B$. The entire state space for the experiment is therefore given by $R \times C \times B^n$, and the whole matrix of choices a is an act in $X^{R \times C \times B^n}$.

We set $n = 20$ throughout our experiment. Probability matching (PM) questions are given by $f = 2_k 0$ and $g = 0_k 2$, where $k \in \{11, 12, \dots, 16\}$. Risky-Safe (RS) questions offer $f = 2_k 0$ and $g = 1$, where again $k \in \{11, 12, \dots, 16\}$. We do not model games here, though the games of strategic certainty (SC) are identical to the PM choices except in framing.

The IND and SIM treatments are as described above. In the CORR treatment only one ball $b_1 \in B$ is drawn, and each a_{ij} pays $a_{ij}(b_1)$. The entire state space is therefore $R \times C \times B$, and so \succeq^* is defined over $X^{R \times C \times B}$. In the SEQ treatment there are n ball draws, as in IND, but now the column chosen for payment (c_i) is drawn in advance, the subject chooses each a_{ij} sequentially, starting at $j = 1$ and proceeding until $j = c_i$. The ONE treatment simply sets $n = 1$.

To model choices, we start by assuming the subject has a preference \succeq^* over the entire choice matrix $a \in X^{R \times C \times B^n}$. This is useful later for describing the assumptions under which our payment mechanism is incentive compatible. But for now our focus is on how the subject chooses across the n replications of a single decision problem. In other words, for each decision problem i , we are interested in studying preferences over $a_i \in X^{C \times B^n}$. To capture this we define \succeq over various a_i by

$$a_i \succeq a'_i \Leftrightarrow \begin{pmatrix} a_i \\ a_i \\ \vdots \\ a_i \end{pmatrix} \succeq^* \begin{pmatrix} a'_i \\ a'_i \\ \vdots \\ a'_i \end{pmatrix}. \quad (2)$$

We can then derive a preference \succeq_0 over single choice objects in X^B by

$$a_{ij} \succeq_0 a'_{ij} \Leftrightarrow (a_{ij}, a_{ij}, \dots, a_{ij}) \succeq (a'_{ij}, a'_{ij}, \dots, a'_{ij}). \quad (3)$$

D.2. Dominance, Monotonicity, and Mixing

Consider a subject who faces only one decision problem D_i , and does so n times. Thus, their only choices are $a_i = (a_{i1}, \dots, a_{in})$. We can view this as equivalent to having m rows but choosing the same vector a_i in every row, because then the draw of the row would be irrelevant.

Given these derived preferences, we can formulate several useful notions of dominance. The first is simply stochastic dominance, while the others are various notions of statewise dominance.

Definition 2. Let ρ be an objective probability measure on (the discrete topology of) B .

1. f *stochastically dominates* g if, for every $x \in X$, $\rho(\{b : f(b) \leq x\}) \leq \rho(\{b : g(b) \leq x\})$.
2. f *B-dominates* g if, for all b , $f(b) \geq g(b)$.
3. a_i *C-dominates* a'_i if, for all j , $a_{ij} \succeq_0 a'_{ij}$.
4. a_i *C-stochastically dominates* a'_i if, for all j , a_{ij} stochastically dominates a'_{ij} .
5. a_i *C \times B^n -dominates* a'_i if, for all j and b_j , $a_{ij}(b_j) \geq a'_{ij}(b_j)$.
6. a *R-dominates* a' if, for all i , $a_i \succeq a'_i$.

In general, an object is said to be *dominant* (under the appropriate notion of dominance) if it dominates all other alternatives. For example, a_i is *C-dominant* if it *C-dominates* every a'_i .³⁵

For each notion of dominance we can also define an equivalent notion of monotonicity (with respect to dominance) of the subject's preference.³⁶

³⁵In Appendix C *C*-dominance was called replicate dominance, and *C*-stochastic dominance was simply called stochastic dominance.

³⁶In earlier drafts *R*-monotonicity was called “row monotonicity” and *C*-monotonicity was called “replicate monotonicity.”

- Definition 3.**
1. \succeq_0 satisfies *stochastic monotonicity* if $f \succeq_0 g$ whenever f stochastically dominates g .
 2. \succeq_0 satisfies *B-monotonicity* if $f \succeq_0 g$ whenever f B-dominates g .
 3. \succeq satisfies *C-monotonicity* if $a_i \succeq a'_i$ whenever a_i C-dominates a'_i .
 4. \succeq satisfies *C-stochastic monotonicity* if $a_i \succeq a'_i$ whenever a_i C-stochastically dominates a'_i .
 5. \succeq satisfies *C × Bⁿ-monotonicity* if $a_i \succeq a'_i$ whenever a_i C × Bⁿ-dominates a'_i .
 6. \succeq^* satisfies *R-monotonicity* if $a \succeq^* a'$ whenever a R-dominates a' .

Each of these can equivalently be defined in terms of deviations in a single state. For example, an equivalent definition of R-monotonicity is that, for all i , a_i , a'_i , and a'' ,

$$\begin{pmatrix} a''_1 \\ \vdots \\ a''_{i-1} \\ a_i \\ a''_{i+1} \\ \vdots \\ a''_m \end{pmatrix} \succeq^* \begin{pmatrix} a'_1 \\ \vdots \\ a''_{i-1} \\ a'_i \\ a''_{i+1} \\ \vdots \\ a''_m \end{pmatrix} \Leftrightarrow a_i \succeq a'_i.$$

And an equivalent definition of C-monotonicity is that, for all i, j , a_{ij} , a'_{ij} , and a'' ,

$$(a''_{i1}, \dots, a''_{ij-1}, a_{ij}, a''_{ij+1}, \dots, a''_{in}) \succeq (a''_{i1}, \dots, a''_{ij-1}, a'_{ij}, a''_{ij+1}, \dots, a''_{in}) \Leftrightarrow a_{ij} \succeq_0 a'_{ij}.^{37}$$

We can also talk about a subject whose preferences satisfy certain monotonicity concepts on some problems, but not others. For example, \succeq may satisfy C-monotonicity on D_i , but not on $D_{i'}$.

In our experiment the main object of focus is \succeq —how people choose across multiple replicates of the same problem. Thus, we want \succeq to be revealed truthfully. Azrieli et al.

³⁷One could instead define C-monotonicity identically to R-monotonicity, mutatis mutandis, by first defining a relation over entire columns and then requiring that this preference be independent of what is chosen in other columns. This would be strictly stronger than our definition of C-monotonicity because ours only applies to the special case where all rows are identical (which corresponds to the case of only having a single decision problem).

(2018) show that this is true if (and, essentially, only if) \succeq^* satisfies R -monotonicity. The argument is simple: Picking the \succeq -most preferred a_i on each i generates matrix a , and any deviation a' would lead to at least one row i on which $a_i > a'_i$. Thus, a R -dominates a' . If \succeq^* satisfies R -dominance, then the subject would never prefer such a deviation. Thus, we assume R -monotonicity throughout, but do not assume any other form of monotonicity listed above. Justification for this comes from Brown and Healy (2018), who show that monotonicity assumptions may be violated when all decisions are shown on the same screen, but not when they are shown on separate screens and in random order. In our experiment the decision problems are shown on separate screens and in random order, so we expect R -monotonicity to hold. The replicates, however, are all shown on the same screen, and so we may expect violations of other forms of monotonicity for \succeq .

R -monotonicity is not innocuous, however. It forces a form of independence across decision problems: if a_i is chosen in row i , then it must be chosen regardless of what was chosen in other rows.³⁸

It is useful to highlight the relationships between the three dominance concepts that apply to \succeq .

- Lemma 1.**
1. \succeq satisfies C -stochastic monotonicity $\Rightarrow \succeq$ satisfies $C \times B^n$ -monotonicity.
 2. Suppose \succeq_0 satisfies B -monotonicity. Then \succeq satisfies C -monotonicity $\Rightarrow \succeq$ satisfies $C \times B^n$ -monotonicity.
 3. Suppose \succeq_0 satisfies stochastic monotonicity. Then \succeq satisfies C -monotonicity $\Rightarrow \succeq$ satisfies C -stochastic monotonicity $\Rightarrow \succeq$ satisfies $C \times B^n$ -monotonicity.

We are interested in studying *mixing* behavior, where subjects vary their choices from one replicate to the next.

³⁸To illustrate, consider a subject facing $D_1 = \{2_90, 1\}$ and $D_2 = \{2_{10}0, 1\}$, each two times (so $m = n = 2$). Suppose his preferences are given by

$$\begin{pmatrix} 1 & 1 \\ 2_{10}0 & 2_{10}0 \end{pmatrix} \succ^* \begin{pmatrix} 1 & 2_90 \\ 1 & 2_{10}0 \end{pmatrix} \succ^* \begin{pmatrix} 1 & 2_90 \\ 2_{10}0 & 2_{10}0 \end{pmatrix} \succ^* \begin{pmatrix} 1 & 1 \\ 1 & 2_{10}0 \end{pmatrix}.$$

This may be because he most-prefers to receive the safe option in exactly two states, but doesn't care which, but does prefer having $2_{10}0$ in place of 2_90 . Unfortunately this violates R -monotonicity since

$$\begin{pmatrix} 1 & 2_90 \\ 1 & 2_{10}0 \end{pmatrix} \succ^* \begin{pmatrix} 1 & 2_90 \\ 2_{10}0 & 2_{10}0 \end{pmatrix} \Rightarrow (1 \ 2_{10}0) \succeq (2_{10}0 \ 2_{10}0) \Rightarrow \begin{pmatrix} 1 & 1 \\ 1 & 2_{10}0 \end{pmatrix} \succeq^* \begin{pmatrix} 1 & 1 \\ 2_{10}0 & 2_{10}0 \end{pmatrix}.$$

Definition 4. A subject exhibits *mixing* on decision problem D_i if there exist replicates j and j' such that $a_{ij} \neq a_{ij'}$.

The various notions of monotonicity of \succeq rule out mixing behavior in different types of problems.

- Proposition 1.**
1. If \succeq satisfies C -monotonicity then the subject will never mix on any decision problem $D_i = \{f_i, g_i\}$, because they will always choose the option (f_i or g_i) that they prefer.
 2. If \succeq satisfies C -stochastic monotonicity then the subject will never mix on any decision problem $D_i = \{f_i, g_i\}$ in which f_i stochastically dominates g_i , because they will always choose f_i .
 3. If \succeq satisfies $C \times B^n$ -monotonicity then the subject will never mix on any decision problem $D_i = \{f_i, g_i\}$ in which f_i B -dominates g_i , because they will always choose f_i .

In our experiment we do not offer decision problems with objects that are ranked by B -dominance; thus, we do not test $C \times B^n$ -monotonicity separately from the other two notions of monotonicity.

As a shorthand, we say that a subject has *convex preferences* if \succeq violates the relevant monotonicity concept. Subjects with convex preferences will exhibit mixing behavior (choosing different options on different replicates) for at least some decision problems.

D.3. Mixing and Random Preferences

An obvious explanation for mixing is that subjects simply have convex preferences, meaning they fail to satisfy C -monotonicity (or C -stochastic monotonicity if the options are ranked by stochastic dominance). An alternative explanation for mixing is that subject's preferences simply change from one choice to the next. We argue that such behavior can persist even when C -monotonicity (appropriately re-interpreted) is satisfied.

To formalize this claim, we adapt the framework of Azrieli et al. (2018, online appendix). Specifically, we model preferences as being affected by some unknown state $\theta \in \Theta$. Information about θ is revealed before each decision is made; to capture this

$\geq_0 \setminus \succeq$	Convex	Linear
Random	RC	RL
Fixed	FC	FL

Table XVII: The general typology of subjects.

simply, we let $\theta = ((\theta_{ij})_{j=1}^n)_{i=1}^m$ and assume that at each decision ij the subject observes $\theta_{ij} \in \Theta_{ij}$.³⁹ The subject selects a *plan* $s = ((s_{ij})_{j=1}^n)_{i=1}^m$, where each $s_{ij} : \Theta_{ij} \rightarrow D_i$ indicates what the subject will pick for every possible θ_{ij} . A plan s therefore generates an act that not only depends on r , c , and b , but also on the realized θ . The preference \succeq^* is now defined over the space of such acts. R -monotonicity and C -monotonicity are defined exactly as above, except now a_{ij} is an act that depends on θ as well as b (it lists what would be chosen for every θ). A plan s^* is *truthful* if, at every ij and θ_{ij} , $s_{ij}^*(\theta_{ij})$ is the most-preferred option in D_{ij} , conditional on observing θ_{ij} . Preferences on a_{ij} are assumed to respect dominance, in the sense that $a_{ij}(\theta_{ij}) \succeq_0 a'_{ij}(\theta_{ij})$ for all θ_{ij} implies $a_{ij} \succeq_0 a'_{ij}$.⁴⁰

In this framework, C -monotonicity guarantees that the subject will report their true favorite choices in each replicate, even as the information they observe about their preferences changes from one replicate to the next (Azrieli et al., 2018). It does *not* guarantee that choices will be identical across replicates, only that they will be truthful. This gives our second explanation for mixing:

Observation 1. A subject with random preferences may mix in some decision problems even if they satisfy C -monotonicity.

Thus we have two general explanations for mixing: random preferences and a failure of C -monotonicity (or C -stochastic monotonicity). For simplicity we say those that satisfy monotonicity have *linear* preferences while those who fail it have convex preferences. We can thus type subjects into four categories, as shown in XVII.

³⁹To capture dynamic information revelation we think of θ_{ij} as including all information from all $\theta_{i'j'}$ for which $i' \leq i$ and $j' \leq j$.

⁴⁰Here, $a_{ij}(\theta_{ij}) \in D_{ij}$ is the constant act that pays the same gamble for all r , c , and θ , and abusing notation, \succeq_0 also represents preferences over these acts.

D.4. Mixing in The Sequential Treatment

We propose instead that the sequential treatment triggers myopic preferences. The idea is that the subject faces a “current choice” and “future choices.” In SEQ the current choice at each j is simple: pick between f_i and g_i . This choice is guided by \succeq_0 over f_i and g_i . The subject ignores future choices. In SIM there is only one “current choice,” which is a choice over the entire vector a_i . This is guided by \succeq .

Formally, let $C(j) \subseteq C$ represent those states in C that are still possible at information set j , but not at information set $j + 1$. In our SEQ treatment, $C(j) = \{j\}$ for all j . In the IND and SIM treatments the only information set is $j = 1$, so $C(1) = C$. For each j , define \succeq^j as the subject’s preference over acts of the form $a_i^j = (a_{ij})_{j \in C(j)} \in X^{C(j) \times B^{\#C(j)}}$.⁴¹

Definition 5. Preference \succeq is *myopic* if, for all information sets j , $a_i^j \succeq^j a_i^{j'}$ then we have $a_i \succeq a_i'$.

This definition does not necessarily pin down the entire ranking \succeq , but it does pin down a most-preferred element. Specifically, if there are J information sets and a_i^j is the most-preferred element at each j according to \succeq^j , then under myopic preferences $a_i = (a_i^1, \dots, a_i^J)$ must be the most-preferred element according to \succeq .

In SEQ $C(j) = \{j\}$ for each j , so $a_i^j = a_{ij}$ and $\succeq^j = \succeq_0$. Having myopic preferences is therefore equivalent to having preferences that respect C -dominance. In IND and SIM, $C(1) = C$, so $a_i^1 = a_i$ and $\succeq^1 = \succeq$. In those treatments myopic preferences place no restriction whatsoever on \succeq ; the definition becomes vacuous.

The SIM treatment occurs after the SEQ treatment. It is possible that subject learn to adapt myopic preferences in the SEQ treatment and apply them in the SIM treatment that follows.

Subjects with random preferences will continue to mix in the SEQ treatment, as \succeq_0 changes from one information set to the next.

⁴¹ $\#C(j)$ denotes the number of states in $C(j)$.

E. EXPERIMENTAL INSTRUCTIONS

The following eight pages reproduce the experimental instructions given to subjects in the IND treatment.

OVERVIEW

Welcome to our experiment. Thank you for participating! Before we begin, please turn off and put away your cell phones, and put away any other items you might have brought with you. If you have any questions during the instruction period, please raise your hand.

This experiment consists of 4 different “blocks.” In each block, you’ll be asked to make a bunch of decisions. (The decisions are numbered, but will appear in random order. For example, you may make decision #7, and then decision #2, and so on.) Your choices in one block will not affect your choices in the other blocks; the four blocks are completely independent. We’ll go over instructions at the start of each block. Your screens will also give instructions, and you’re free to refer back to the printed instructions at any time.

At the end of the experiment, one of the decisions will be randomly selected for payment. In each decision we will describe how that decision gets paid if it is selected.

In addition to being paid for one decision, you will also receive a \$5 participation payment for completing the experiment.

BINGO CAGE BETS

We have a Bingo cage filled with 20 balls, numbered 1-20.

In each question in this block, you will be offered two “bets” on which ball is drawn from the cage. We’ll actually draw a ball from the Bingo cage 20 times, and you’ll choose 20 bets, one for each draw. (After each draw we’ll put the ball back into the cage before the next draw.) In each decision you must choose between Bet A or Bet B, both of which will be shown on your computer screen. Here is an example of two bets you might be given:

Bet A: You receive \$15 regardless of which ball is drawn.	Bet B: You receive \$25 if the ball drawn is from 1–16, and \$5 if it is from 17–20.																																																												
<p>\$15</p> <table style="margin: auto;"><tr><td>1</td><td>2</td><td>3</td><td>4</td></tr><tr><td>5</td><td>6</td><td>7</td><td>8</td></tr><tr><td>9</td><td>10</td><td>11</td><td>12</td></tr><tr><td>13</td><td>14</td><td>15</td><td>16</td></tr><tr><td>17</td><td>18</td><td>19</td><td>20</td></tr></table>	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	<table style="margin: auto;"><tr><td colspan="4">\$25</td><td colspan="4">\$5</td></tr><tr><td>1</td><td>2</td><td>3</td><td>4</td><td>17</td><td>18</td><td>19</td><td>20</td></tr><tr><td>5</td><td>6</td><td>7</td><td>8</td><td></td><td></td><td></td><td></td></tr><tr><td>9</td><td>10</td><td>11</td><td>12</td><td></td><td></td><td></td><td></td></tr><tr><td>13</td><td>14</td><td>15</td><td>16</td><td></td><td></td><td></td><td></td></tr></table>	\$25				\$5				1	2	3	4	17	18	19	20	5	6	7	8					9	10	11	12					13	14	15	16				
1	2	3	4																																																										
5	6	7	8																																																										
9	10	11	12																																																										
13	14	15	16																																																										
17	18	19	20																																																										
\$25				\$5																																																									
1	2	3	4	17	18	19	20																																																						
5	6	7	8																																																										
9	10	11	12																																																										
13	14	15	16																																																										

Payment:

If one of these questions is chosen for payment, we’ll draw a ball from the Bingo cage 20 times. We’ll then roll a 20-sided die to determine which of the 20 draws to pay out. We’ll then look at which bet you chose for that draw, and pay you based on that draw.

For example, suppose the 20-sided die roll comes up “3”. That means we’re paying you for the bet you chose on the 3rd draw of the ball. Suppose you chose Bet B, shown above. Bet B pays \$25 if the ball is 1-16.

If the 20 draws from the Bingo cage are

5, 3, 11, 5, 20, 8, 4, 9, 1, 15, 9, 9, 11, 2, 18, 12, 5, 8, 12, 10

then the 3rd draw is 11. You chose Bet B, and Bet B pays \$25 for ball 11, so you’d actually be paid \$25.

If the 20-sided die had come up “5” then we’d pay for the 5th draw, which is 20. In that case Bet B would only pay \$5.

If you had chosen Bet A then you’d receive \$15 regardless of what ball is drawn.

The actual bets offered may be different than this example, and you’ll make several choices like this. Read the description of the 2 bets carefully each time before making your 20 choices.

GAMES AGAINST PAST PLAYERS

In these questions, you will play a “matrix game” against 20 people who participated in this experiment on some prior date.

On the screen we will now demonstrate how “matrix games” work.

In this block, you will be the ROW player and the past participants were COLUMN players.

ROW player choices:

You will actually play each game 20 times. For each of your 20 choices we will randomly draw one of the 20 past participants, and your choice will be paired against that past participant’s choice. But you won’t know which past participants you’re paired with in each choice until the end of the experiment.

Before you make your 20 decisions, we might give you some information about what all 20 past participants chose. For example, we could tell you that of the 20 past participants, 12 chose Left and 8 chose Right. This information will appear on your computer screen.

Payment:

If one of these games is chosen for payment, we’ll use draws from a Bingo cage to see which past participant is associated to each of your 20 choices (putting the ball back after each draw), and then we’ll roll a 20-sided die to see which of those choices is paid out. We’ll compare your Row choice to that person’s Column choice and pay you your payoff in the game for that Row and Column. (The Column player will not be paid; they were paid when they played this game previously.)

GAMES AGAINST CURRENT PLAYERS

In these questions, you will play a “matrix game” against one of the 20 other people in the room today.

On the screen we will now demonstrate how “matrix games” work.

In this block, you will play each game as the Column player *and* as the Row player. You’ll actually proceed through 5 “Stages” of decision-making, numbered Stage 0 through Stage 4. We’ll explain each now:

“Stage 0:” COLUMN player choice:

In Stage 0 you will play the game 1 time as the COLUMN player. Here is an example game:

	Left	Right
Up	\$25 \$5	\$5 \$25
Down	\$5 \$25	\$25 \$5

I choose

“Stage 1:” ROW player choices:

In Stage 1 you now play the same game, but as the ROW player. And you’ll play it 20 times. For each choice we’ll randomly draw the ID of another person in your room, and they will serve as the Column player if that choice is chosen for payment. For example, if your 3rd choice is against Column Player #17, then your 3rd Row choice will be compared to Player #17’s Column choice from Stage 0.

Here is an example screenshot of your 20 choices:

	Left	Right
Up	\$25 \$5	\$5 \$25
Down	\$5 \$25	\$25 \$5

For choice number 1 , I choose	<input type="button" value="Down"/>	For choice number 2 , I choose	<input type="button" value="Up"/>	For choice number 3 , I choose	<input type="button" value="Up"/>
For choice number 4 , I choose	<input type="button" value="Down"/>	For choice number 5 , I choose	<input type="button" value="Up"/>	For choice number 6 , I choose	<input type="button" value="Up"/>
For choice number 7 , I choose	<input type="button" value="Up"/>	For choice number 8 , I choose	<input type="button" value="Down"/>	For choice number 9 , I choose	<input type="button" value="Down"/>
For choice number 10 , I choose	<input type="button" value="Up"/>	For choice number 11 , I choose	<input type="button" value="Down"/>	For choice number 12 , I choose	<input type="button" value="Down"/>
For choice number 13 , I choose	<input type="button" value="Up"/>	For choice number 14 , I choose	<input type="button" value="Up"/>	For choice number 15 , I choose	<input type="button" value="Up"/>
For choice number 16 , I choose	<input type="button" value="Down"/>	For choice number 17 , I choose	<input type="button" value="Up"/>	For choice number 18 , I choose	<input type="button" value="Up"/>
For choice number 19 , I choose	<input type="button" value="Down"/>	For choice number 20 , I choose	<input type="button" value="Down"/>		

Payment:

If Stage 1 is chosen for payment, we'll randomly select one person in the room to be our Row player. And then we'll use draws from a Bingo cage to select the identity of the Column player for each of their 20 choices (putting the ball back after each draw). Finally, we'll use a 20-sided die to see which choice is paid out. That Row player and Column player will get paid based on how they played (the Row player is paid for their Row choice against that particular Column player, and the Column player is paid based on their Column choice from Stage 0.)

Everyone else in the room will receive a fixed payment of \$15.

“Stage 2:” Probabilities:

In Stage 2 we want to know how likely you think it is that Column players play “Left” in this game. One way we could do this is to ask you the following list of 100 questions:

Q#		Option A		Option B
1	Would you rather have	\$20 if COLUMN chose Left	or	1% chance of \$20
2	Would you rather have	\$20 if COLUMN chose Left	or	2% chance of \$20
3	Would you rather have	\$20 if COLUMN chose Left	or	3% chance of \$20
⋮	⋮	⋮	⋮	⋮
99	Would you rather have	\$20 if COLUMN chose Left	or	99% chance of \$20
100	Would you rather have	\$20 if COLUMN chose Left	or	100% chance of \$20

In each question, you'd pick either Option A or Option B. Presumably you'd want Option A in the first few questions, but at some point would switch to taking Option B. So rather than telling us your choice to all 100 questions, we can just ask you to tell us at what percent chance you'd switch. And that “switch point” is exactly where you're indifferent between Option A and Option B, because that switch point would be exactly at the probability that you think the Column players are choosing Left.

For example, suppose your switch point is 73%. That means you're indifferent between getting \$20 if COLUMN plays Left, and getting \$20 with 73% chance. But if you're indifferent between those choices, then you must think COLUMN is playing Left 73% of the time. In other words, your switch point is exactly your probability that they play Left.

How would you be paid if Stage 2 is chosen for payment? You enter your probability that the Column player plays left (for example, 73%). Then we draw one of the 100 questions above and see what you'd choose on that question. If it's #1-72 then you chose Option A. So we'd pay you \$20 if a randomly-selected Column player actually chose Left in Stage 0. If the question drawn is #73-100 then you chose Option B. So we'd pay you \$20 with the probability given in that row. (For example, if we pick question #83, then you'd get \$20 with an 83% chance.)

We'll use two 10-sided dice to pick which row is actually chosen. If you choose Option B then we'll use another roll of the two 10-sided dice to determine whether you win the \$20 or not. (For example, if the chosen row is #83, then you're getting an 83% chance of \$20. That means we'll pay you \$20 if the second roll comes up 1-83.)

Obviously you have an incentive to announce your "true" probability that you think the Column player is playing Left. If you misreport your true probability then you'll end up choosing an option you like less on some of the rows above.

Here is an example screenshot of this decision:

	Left	Right
Up	\$25 \$5	\$5 \$25
Down	\$5 \$25	\$25 \$5

I think the probability that they chose Left is %.

"Stage 3:" Row player with a Hint

In Stage 3 we'll show you a "hint" of how an actual Column player played today. Here's how the hint works:

First, the computer will randomly select 1 of the other 20 players. The computer knows whether this player chose Left or Right as COLUMN player, so the computer can give you a hint about which they chose. The hint will either say "Left" or "Right", but it's not very accurate; the hint will be correct 55% of the time and wrong 45% of the time.

This means that if you see the hint that COLUMN chose Left, then it's slightly more likely that the COLUMN player really did choose Left. And if the hint says "Right" then it's slightly more likely that the COLUMN player really did choose Right.

After you see this hint, you will play the game 20 more times as the ROW player, each time matched with a randomly-drawn person in the room, just as you did back in Stage 1. The only difference is that you've now seen a hint.

"Stage 4:" Probabilities with a Hint

In Stage 4 we'll once again ask you your probability that a random Column player chose Left. The payments will work just like in Stage 2. Again, your incentive is to report your belief truthfully. This is exactly as in Stage 2; the only difference is that now you've seen a hint.

You will play 2 matrix games in this block. In each game you will go through all 5 stages (0 through 4). Notice that the games' payoffs may be different, but the procedures for each stage are exactly the same.

INVESTMENT QUESTIONS

In the investment questions, you will be given \$10.00, and you can choose to invest any amount of that money in a risky project. The money you don't invest you keep for yourself.

The project can either be a success or a failure.

If it's successful then the amount you invested in it will be multiplied by some number and paid to you. If it's a failure then that money will be lost.

In either case you get to keep the money that you chose *not* to invest.

Your screen will include detailed instructions about these questions, so read the information carefully. Here is an example screenshot:

- The risky project has a **40%** chance of succeeding.
- If it succeeds, the money you invested will be multiplied by **3**.
- If it does not succeed, the money you invested is lost.
- You always keep any money that you did not invest.

I choose to invest \$ of my \$10.00 in the risky project.
(The remaining \$2.77 I will not invest.)

You will face two different investment choices, each with a different chance of success and multiplier. Please read the screen carefully before making your choice each time.

The following eight pages reproduce the COR treatment instructions.

OVERVIEW

Welcome to our experiment. Thank you for participating! Before we begin, please turn off and put away your cell phones, and put away any other items you might have brought with you. If you have any questions during the instruction period, please raise your hand.

This experiment consists of 4 different “blocks.” In each block, you’ll be asked to make a bunch of decisions. (The decisions are numbered, but will appear in random order. For example, you may make decision #7, and then decision #2, and so on.) Your choices in one block will not affect your choices in the other blocks; the four blocks are completely independent. We’ll go over instructions at the start of each block. Your screens will also give instructions, and you’re free to refer back to the printed instructions at any time.

At the end of the experiment, one of the decisions will be randomly selected for payment. In each decision we will describe how that decision gets paid if it is selected.

In addition to being paid for one decision, you will also receive a \$5 participation payment for completing the experiment.

BINGO CAGE BETS

We have a Bingo cage filled with 20 balls, numbered 1-20.

In each question in this block, you will be offered two “bets” on which ball is drawn from the cage. We’ll only draw a ball from the Bingo cage 1 time, but you will make 20 bets on that one draw. In each you must choose between Bet A or Bet B, both of which will be shown on your computer screen. Here is an example of two bets you might be given:

Bet A: You receive \$15 regardless of which ball is drawn.	Bet B: You receive \$25 if the ball drawn is from 1–16, and \$5 if it is from 17–20.																																																												
<p>\$15</p> <table style="margin: auto;"><tr><td>1</td><td>2</td><td>3</td><td>4</td></tr><tr><td>5</td><td>6</td><td>7</td><td>8</td></tr><tr><td>9</td><td>10</td><td>11</td><td>12</td></tr><tr><td>13</td><td>14</td><td>15</td><td>16</td></tr><tr><td>17</td><td>18</td><td>19</td><td>20</td></tr></table>	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	<table style="margin: auto;"><tr><td colspan="4">\$25</td><td colspan="4">\$5</td></tr><tr><td>1</td><td>2</td><td>3</td><td>4</td><td>17</td><td>18</td><td>19</td><td>20</td></tr><tr><td>5</td><td>6</td><td>7</td><td>8</td><td></td><td></td><td></td><td></td></tr><tr><td>9</td><td>10</td><td>11</td><td>12</td><td></td><td></td><td></td><td></td></tr><tr><td>13</td><td>14</td><td>15</td><td>16</td><td></td><td></td><td></td><td></td></tr></table>	\$25				\$5				1	2	3	4	17	18	19	20	5	6	7	8					9	10	11	12					13	14	15	16				
1	2	3	4																																																										
5	6	7	8																																																										
9	10	11	12																																																										
13	14	15	16																																																										
17	18	19	20																																																										
\$25				\$5																																																									
1	2	3	4	17	18	19	20																																																						
5	6	7	8																																																										
9	10	11	12																																																										
13	14	15	16																																																										

Payment:

If one of these questions is chosen for payment, we’ll draw a ball from the Bingo cage. We’ll then roll a 20-sided die to determine which of the 20 bets you chose will be paid.

For example, suppose the 20-sided die roll comes up “3”. That means we’re paying you for the 3rd choice you made. Suppose on the 3rd choice you chose Bet B, shown above. Bet B pays \$25 if the ball is 1-16.

Suppose the ball drawn from the Bingo cage were ball #11. You chose Bet B, and Bet B pays \$25 for ball 11, so you’d actually be paid \$25.

If the 20-sided die had come up “5” then we’d pay for your 5th chosen bet. If that’s also Bet B then you’d also get \$25. But if your 5th chosen bet were Bet A then you’d get \$15.

The actual bets offered may be different than this example, and you’ll make several choices like this. Read the description of the 2 bets carefully each time before making your 20 choices.

GAMES AGAINST PAST PLAYERS

In these questions, you will play a “matrix game” against someone who participated in this experiment several weeks ago.

On the screen we will now demonstrate how “matrix games” work.

In this block, you will be the ROW player and the past participants were COLUMN players.

ROW player choices:

You will actually play each game 20 times against one past participant. That past participant is one of 20 people who played previously.

Before you make your 20 decisions, we might give you some information about what the whole group of past participants chose. For example, we could tell you that of the 20 past participants, 12 chose Left and 8 chose Right. This information will appear on your computer screen.

Payment:

If one of these games is chosen for payment, we’ll roll a 20-sided die to see which of your Row choices counts for payment. We’ll then compare your Row choice to that past Column player’s choice, and pay you your payoff in the game for that Row and Column. (The Column player will not be paid; they were paid when they played this game several weeks ago.)

GAMES AGAINST CURRENT PLAYERS

In these questions, you will play a “matrix game” against one other person in the room today.

On the screen we will now demonstrate how “matrix games” work.

In this block, you will play each game as the Column player *and* as the Row player. You’ll actually proceed through 5 “Stages” of decision-making, numbered Stage 0 through Stage 4. We’ll explain each now:

“Stage 0:” COLUMN player choice:

In Stage 0 you will play the game 1 time as the COLUMN player. Here is an example game:

	Left	Right
Up	\$25 \$5	\$5 \$25
Down	\$5 \$25	\$25 \$5

I choose
Left
Right

“Stage 1:” ROW player choices:

In Stage 1 you now play the same game, but as the ROW player. You’ll be matched with one Column player, but you’ll choose how to play the game against that player 20 different times. Only one of your 20 choices will be selected randomly for payment.

Here is an example screenshot of these 20 choices:

	Left	Right
Up	\$25 \$5	\$5 \$25
Down	\$5 \$25	\$25 \$5

We will randomly select one player in the room.
You will play this game 20 times, each against this same player.
You play against his or her Column choice from the previous screen.
The plays of the game are numbered from 1 to 20 below.

For play number 1, I choose ▾ For play number 2, I choose ▾ For play number 3, I choose ▾
For play number 4, I choose ▾ For play number 5, I choose ▾ For play number 6, I choose ▾
For play number 7, I choose ▾ For play number 8, I choose ▾ For play number 9, I choose ▾
For play number 10, I choose ▾ For play number 11, I choose ▾ For play number 12, I choose ▾
For play number 13, I choose ▾ For play number 14, I choose ▾ For play number 15, I choose ▾
For play number 16, I choose ▾ For play number 17, I choose ▾ For play number 18, I choose ▾
For play number 19, I choose ▾ For play number 20, I choose ▾

Payment:

If Stage 1 is chosen for payment, we’ll randomly select one person in the room to be our Row player. Then we’ll randomly select one other person to be the Column. Finally, we’ll randomly select which of the 20 Row player choices will be played for real. Those two players will get paid based on how they played (the Row players is paid for their randomly-selected Row choice, and the Column player is paid based on their Column choice from Stage 0.)

Everyone else in the room will receive a fixed payment of \$15.

“Stage 2:” Probabilities:

In Stage 2 we want to know how likely you think it is that Column players play “Left” in this game. One way we could do this is to ask you the following list of 100 questions:

Q#		Option A		Option B
1	Would you rather have	\$20 if COLUMN chose Left	or	1% chance of \$20
2	Would you rather have	\$20 if COLUMN chose Left	or	2% chance of \$20
3	Would you rather have	\$20 if COLUMN chose Left	or	3% chance of \$20
⋮	⋮	⋮	⋮	⋮
99	Would you rather have	\$20 if COLUMN chose Left	or	99% chance of \$20
100	Would you rather have	\$20 if COLUMN chose Left	or	100% chance of \$20

In each question, you’d pick either Option A or Option B. Presumably you’d want Option A in the first few questions, but at some point would switch to taking Option B. So rather than telling us your choice to all 100 questions, we can just ask you to tell us at what percent chance you’d switch. And that “switch point” is exactly where you’re indifferent between Option A and Option B, because that switch point would be exactly at the probability that you think the Column players are choosing Left.

For example, suppose your switch point is 73%. That means you’re indifferent between getting \$20 if COLUMN plays Left, and getting \$20 with 73% chance. But if you’re indifferent between those choices, then you must think COLUMN is playing Left 73% of the time. In other words, your switch point is exactly your probability that they play Left.

How would you be paid if Stage 2 is chosen for payment? You enter your probability that the Column player plays left (for example, 73%). Then we draw one of the 100 questions above and see what you’d choose on that question. If it’s #1-72 then you chose Option A. So we’d pay you \$20 if a randomly-selected Column player actually chose Left in Stage 0. If the question drawn is #73-100 then you chose Option B. So we’d pay you \$20 with the probability given in that row. (For example, if we pick question #83, then you’d get \$20 with an 83% chance.)

We’ll use two 10-sided dice to pick which row is actually chosen. If you choose Option B then we’ll use another roll of the two 10-sided dice to determine whether you win the \$20 or not. (For example, if the chosen row is #83, then you’re getting an 83% chance of \$20. That means we’ll pay you \$20 if the second roll comes up 1-83.)

Obviously you have an incentive to announce your “true” probability that you think the Column player is playing Left. If you misreport your true probability then you’ll end up choosing an option you like less on some of the rows above.

Here is an example screenshot of this decision:

	Left	Right
Up	\$25 \$5	\$5 \$25
Down	\$5 \$25	\$25 \$5

I think the probability that they chose Left is %.

“Stage 3:” Row player with a Hint

In Stage 3 we’ll show you a “hint” of how an actual Column player played today. Here’s how the hint works:

First, the computer will randomly select 1 of the other 20 players. The computer knows whether this player chose Left or Right as COLUMN player, so the computer can give you a hint about which they chose. The hint will either say “Left” or “Right”, but it’s not very accurate; the hint will be correct 55% of the time and wrong 45% of the time.

This means that if you see the hint that COLUMN chose Left, then it’s slightly more likely that the COLUMN player really did choose Left. And if the hint says “Right” then it’s slightly more likely that the COLUMN player really did choose Right.

After you see this hint, you will play the game 20 more times as the ROW player, all against the same Column player, just as you did back in Stage 1. The only difference is that you’ve now seen a hint.

“Stage 4:” Probabilities with a Hint

In Stage 4 we’ll once again ask you your probability that a random Column player chose Left. The payments will work just like in Stage 2. Again, your incentive is to report your belief truthfully. This is exactly as in Stage 2; the only difference is that now you’ve seen a hint.

You will play 2 matrix games in this block. In each game you will go through all 5 stages (0 through 4). Notice that the games’ payoffs may be different, but the procedures for each stage are exactly the same.

INVESTMENT QUESTIONS

In the investment questions, you will be given \$10.00, and you can choose to invest any amount of that money in a risky project. The money you don't invest you keep for yourself.

The project can either be a success or a failure.

If it's successful then the amount you invested in it will be multiplied by some number and paid to you. If it's a failure then that money will be lost.

In either case you get to keep the money that you chose *not* to invest.

Your screen will include detailed instructions about these questions, so read the information carefully. Here is an example screenshot:

- The risky project has a **40%** chance of succeeding.
- If it succeeds, the money you invested will be multiplied by **3**.
- If it does not succeed, the money you invested is lost.
- You always keep any money that you did not invest.

I choose to invest \$ of my \$10.00 in the risky project.
(The remaining \$2.77 I will not invest.)

You will face two different investment choices, each with a different chance of success and multiplier. Please read the screen carefully before making your choice each time.

The following five pages reproduce the SEQ experiment instructions.

OVERVIEW

Welcome to our experiment. Thank you for participating! Before we begin, please turn off and put away your cell phones, and put away any other items you might have brought with you. If you have any questions during the instruction period, please raise your hand.

This experiment consists of 4 different “blocks.” In each block, you’ll be asked to make a bunch of decisions. (The decisions are numbered, but will appear in random order. For example, you may make decision #7, and then decision #2, and so on.) Your choices in one block will not affect your choices in the other blocks; the four blocks are completely independent. We’ll go over instructions at the start of each block. Your screens will also give instructions, and you’re free to refer back to the printed instructions at any time.

At the end of the experiment, one of the decisions will be randomly selected for payment. In each decision we will describe how that decision gets paid if it is selected.

In addition to being paid for one decision, you will also receive a \$5 participation payment for completing the experiment.

BINGO CAGE BETS

We have a Bingo cage filled with 20 balls, numbered 1-20.

In each question in this block, you will be offered two “bets” on which ball is drawn from the cage. These are labeled Bet A and Bet B. You’ll choose between Bet A and Bet B, but you’ll do this up to 20 times. For each of your choices, we’ll draw a ball from the Bingo cage to see what you would be paid. (After each draw we’ll put the ball back into the cage before the next draw.) Here is an example of two bets you might be given:

Bet A: You receive \$15 regardless of which ball is drawn.	Bet B: You receive \$25 if the ball drawn is from 1–16, and \$5 if it is from 17–20.																																																												
<div style="margin-bottom: 10px;"><u> </u> \$15</div> <table style="margin: auto;"><tr><td>1</td><td>2</td><td>3</td><td>4</td></tr><tr><td>5</td><td>6</td><td>7</td><td>8</td></tr><tr><td>9</td><td>10</td><td>11</td><td>12</td></tr><tr><td>13</td><td>14</td><td>15</td><td>16</td></tr><tr><td>17</td><td>18</td><td>19</td><td>20</td></tr></table>	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	<table style="margin: auto;"><tr><td colspan="4"><u> </u> \$25</td><td colspan="4"><u> </u> \$5</td></tr><tr><td>1</td><td>2</td><td>3</td><td>4</td><td>17</td><td>18</td><td>19</td><td>20</td></tr><tr><td>5</td><td>6</td><td>7</td><td>8</td><td></td><td></td><td></td><td></td></tr><tr><td>9</td><td>10</td><td>11</td><td>12</td><td></td><td></td><td></td><td></td></tr><tr><td>13</td><td>14</td><td>15</td><td>16</td><td></td><td></td><td></td><td></td></tr></table>	<u> </u> \$25				<u> </u> \$5				1	2	3	4	17	18	19	20	5	6	7	8					9	10	11	12					13	14	15	16				
1	2	3	4																																																										
5	6	7	8																																																										
9	10	11	12																																																										
13	14	15	16																																																										
17	18	19	20																																																										
<u> </u> \$25				<u> </u> \$5																																																									
1	2	3	4	17	18	19	20																																																						
5	6	7	8																																																										
9	10	11	12																																																										
13	14	15	16																																																										

Payment:

The computer will randomly pick how many times you’ll be asked to choose between Bet A and Bet B. This could be any number from 1 to 20, with each number being equally likely. But you won’t know how many times you get to choose; instead, you’ll make your first choice between A and B and find out if it’s paid. If it’s not, then you’ll make your second choice between A and B and find out if it’s paid. If it’s not, you’ll continue to your third choice between A and B, and so on. You’ll continue making choices until you find out which is paid. Once you find out which choice is paid, you won’t make any more choices on that screen.

If one of these questions is chosen for payment, we’ll draw a ball from the Bingo cage for each of your choices and you’ll be paid based on the *last* draw.

For example, suppose you got to make 12 choices. That means we’ll draw 12 balls from the Bingo cage. Suppose on your 12th (and last) choice you chose Bet B, shown above. Bet B pays \$25 if the ball is 1-16.

If the 12 draws from the Bingo cage are

5, 3, 11, 5, 20, 8, 4, 9, 1, 15, 9, 11

then the 12th draw is 11. You chose Bet B, and Bet B pays \$25 for ball 11, so you’d actually be paid \$25.

As another example, suppose you got to make 5 choices, and the 5 draws from the Bingo cage are

4, 19, 4, 8, 20.

The 5th draw is 20. If you chose Bet B on your 5th (and final) draw then you’d get paid \$5. If you had chosen Bet A then you’d receive \$15.

The actual bets offered may be different than this example, and you’ll make several choices like this. Read the description of the 2 bets carefully each time before making your 20 choices.

BINGO CAGE BETS: 20 SIMULTANEOUS CHOICES

This block is identical to the previous block, except now you will make *all 20* choices.

On each screen you will be offered a choice between Bet A and Bet B, just like before. But now you will make all 20 possible choices between Bet A and Bet B. You'll submit all 20 of your choices, and at the end of the experiment we'll roll a 20-sided die to see which of your 20 choices is actually paid.

Again, an example of a choice you might face is shown here:

Bet A: You receive \$15 regardless of which ball is drawn. <div style="text-align: center; margin-top: 10px;">\$15 <hr style="width: 50%; margin: 0 auto;"/><table style="margin: 0 auto; text-align: center;"><tr><td>1</td><td>2</td><td>3</td><td>4</td></tr><tr><td>5</td><td>6</td><td>7</td><td>8</td></tr><tr><td>9</td><td>10</td><td>11</td><td>12</td></tr><tr><td>13</td><td>14</td><td>15</td><td>16</td></tr><tr><td>17</td><td>18</td><td>19</td><td>20</td></tr></table></div>	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	Bet B: You receive \$25 if the ball drawn is from 1–16, and \$5 if it is from 17–20. <div style="display: flex; justify-content: space-around; margin-top: 10px;"><div style="text-align: center;">\$25 <hr style="width: 50%; margin: 0 auto;"/><table style="margin: 0 auto; text-align: center;"><tr><td>1</td><td>2</td><td>3</td><td>4</td></tr><tr><td>5</td><td>6</td><td>7</td><td>8</td></tr><tr><td>9</td><td>10</td><td>11</td><td>12</td></tr><tr><td>13</td><td>14</td><td>15</td><td>16</td></tr></table></div><div style="text-align: center;">\$5 <hr style="width: 50%; margin: 0 auto;"/><table style="margin: 0 auto; text-align: center;"><tr><td>17</td><td>18</td><td>19</td><td>20</td></tr></table></div></div>	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
1	2	3	4																																						
5	6	7	8																																						
9	10	11	12																																						
13	14	15	16																																						
17	18	19	20																																						
1	2	3	4																																						
5	6	7	8																																						
9	10	11	12																																						
13	14	15	16																																						
17	18	19	20																																						

Payment:

If one of these questions is chosen for payment, we'll draw a ball from the Bingo cage 20 times, one for each of your 20 choices. We'll then roll a 20-sided die to determine which of the 20 choices to pay out. We'll then look at which bet you chose for that draw, and pay you based on that draw.

For example, suppose the 20-sided die roll comes up "3". That means we're paying you for the bet you chose on the 3rd draw of the ball. Suppose you chose Bet B, shown above. Bet B pays \$25 if the ball is 1-16.

If the 20 draws from the Bingo cage are

5, 3, 11, 5, 20, 8, 4, 9, 1, 15, 9, 9, 11, 2, 18, 12, 5, 8, 12, 10

then the 3rd draw is 11. You chose Bet B, and Bet B pays \$25 for ball 11, so you'd actually be paid \$25.

If the 20-sided die had come up "5" then we'd pay for the 5th draw, which is 20. In that case Bet B would only pay \$5.

If you had chosen Bet A then you'd receive \$15 regardless of what ball is drawn.

The actual bets offered may be different than this example, and you'll make several choices like this. Read the description of the 2 bets carefully each time before making your 20 choices.

BINGO CAGE BETS: ONE CHOICE ONLY

This block is identical to the previous block, except now you will make *only one choice*.

On each screen you will be offered a choice between Bet A and Bet B, just like before. But now you will make only one choice between Bet A and Bet B. If that screen is chosen for payment, then that one choice will be paid.

Again, an example of a choice you might face is shown here:

Bet A: You receive \$15 regardless of which ball is drawn.	Bet B: You receive \$25 if the ball drawn is from 1–16, and \$5 if it is from 17–20.																																																												
<p>\$15</p> <table><tr><td>1</td><td>2</td><td>3</td><td>4</td></tr><tr><td>5</td><td>6</td><td>7</td><td>8</td></tr><tr><td>9</td><td>10</td><td>11</td><td>12</td></tr><tr><td>13</td><td>14</td><td>15</td><td>16</td></tr><tr><td>17</td><td>18</td><td>19</td><td>20</td></tr></table>	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	<table><tr><td colspan="4">\$25</td><td colspan="4">\$5</td></tr><tr><td>1</td><td>2</td><td>3</td><td>4</td><td>17</td><td>18</td><td>19</td><td>20</td></tr><tr><td>5</td><td>6</td><td>7</td><td>8</td><td></td><td></td><td></td><td></td></tr><tr><td>9</td><td>10</td><td>11</td><td>12</td><td></td><td></td><td></td><td></td></tr><tr><td>13</td><td>14</td><td>15</td><td>16</td><td></td><td></td><td></td><td></td></tr></table>	\$25				\$5				1	2	3	4	17	18	19	20	5	6	7	8					9	10	11	12					13	14	15	16				
1	2	3	4																																																										
5	6	7	8																																																										
9	10	11	12																																																										
13	14	15	16																																																										
17	18	19	20																																																										
\$25				\$5																																																									
1	2	3	4	17	18	19	20																																																						
5	6	7	8																																																										
9	10	11	12																																																										
13	14	15	16																																																										

Payment:

If one of these screens is chosen for payment, we'll draw a single ball from the Bingo cage. We'll then look at which bet you chose, and pay you based on that one draw.

For example, suppose you chose Bet B, shown above. Bet B pays \$25 if the ball is 1-16. Suppose we draw ball #11. You chose Bet B, and Bet B pays \$25 for ball 11, so you'd actually be paid \$25.

If you had chosen Bet A then you'd receive \$15 regardless of what ball is drawn.

The actual bets offered may be different than this example, and you'll make several choices like this. Read the description of the 2 bets carefully each time before making your choices.

INVESTMENT QUESTIONS

In the investment questions, you will be given \$10.00, and you can choose to invest any amount of that money in a risky project. The money you don't invest you keep for yourself.

The project can either be a success or a failure.

If it's successful then the amount you invested in it will be multiplied by some number and paid to you. If it's a failure then that money will be lost.

In either case you get to keep the money that you chose *not* to invest.

Your screen will include detailed instructions about these questions, so read the information carefully. Here is an example screenshot:

- The risky project has a **40%** chance of succeeding.
- If it succeeds, the money you invested will be multiplied by **3**.
- If it does not succeed, the money you invested is lost.
- You always keep any money that you did not invest.

I choose to invest \$ of my \$10.00 in the risky project.
(The remaining \$2.77 I will not invest.)

You will face two different investment choices, each with a different chance of success and multiplier. Please read the screen carefully before making your choice each time.

The following six pages reproduce the Negative-Demand online experiment. The online replication of IND had the exact same instructions, except the sentence “You would be doing us a favor if you chose the same bet in all 20 choices.” is removed from the fourth page.



Caltech

Thank you for your participation!

If you complete this study, you will receive a \$3 completion payment. One out of every 10 participants will be randomly selected to receive an additional bonus payment based on their decisions throughout the study.

Please note that all payments in this experiment are in US Dollars (\$). If you are using a different local currency, it will be converted from US Dollars at the current exchange rate.

→

Powered by Qualtrics [↗](#)



Caltech

There are 4 question groups in this study. In each question group, you will make 20 choices. In each choice, you will choose which of two bets you prefer. The bets specify the amount of bonus payment you would receive if that decision were randomly chosen to determine your bonus payment. In each bet, the actual bonus payment you would receive depends on a random number from 1-20 that will be drawn, with each number equally likely.

Each of your 20 bets will correspond to a different random number drawn. For each of your 20 bets, the computer will draw a random number 1-20, each number equally likely. The outcome of your 20 bets will be determined by each of these 20 random numbers.

In each choice you must choose between Bet A or Bet B, both of which will be shown on your computer screen. Here is an example of two bets you might be given:

Bet A: You receive \$25 if the number drawn is from 1--11, and \$5 if it is from 12--20.	Bet B: You receive \$5 if the number drawn is from 1--11, and \$25 if it is from 12--20.
<div style="display: flex; justify-content: space-around;"> <div style="text-align: center;"> <p>\$25</p> </div> <div style="text-align: center;"> <p>\$5</p> </div> </div>	<div style="display: flex; justify-content: space-around;"> <div style="text-align: center;"> <p>\$5</p> </div> <div style="text-align: center;"> <p>\$25</p> </div> </div>

Choice 1	Choice 2	Choice 3	Choice 4	Choice 5
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>
Choice 6	Choice 7	Choice 8	Choice 9	Choice 10
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>
Choice 11	Choice 12	Choice 13	Choice 14	Choice 15
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>
Choice 16	Choice 17	Choice 18	Choice 19	Choice 20
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>

If you are randomly selected to receive a bonus payment, we'll draw 20 random numbers 1-20. These random numbers will determine the outcome from each of your 20 bets.

Then, we'll draw another number 1-20 to determine which of your choices to pay out. We'll look at which bet you chose for that choice, and pay you based on the outcome as determined by the random number corresponding to that choice.

For example, suppose we draw the following 20 random numbers: 5, 3, 11, 5, 20, 8, 4, 9, 1, 15, 9, 9, 11, 2, 18, 12, 5, 8, 12, 10. This means the random number drawn for Choice 1 is 5, the random number drawn for Choice 2 is 3, the random number drawn for Choice 3 is 11, etc.

Then we would randomly draw another number 1-20, each equally likely, to determine which of your choices is the one that will determine your bonus payment. Let's say that this number is 3. This means that you would be paid according to whichever bet you chose in Choice 3. And, as

specified above, the randomly selected number for choice 3 is 11. If you chose Bet A in Choice 3, you'd receive a \$25 bonus payment (since Bet A pays \$25 for numbers 1-11). If you chose Bet B in Choice 3, you'd receive a \$5 bonus payment (since Bet B pays \$5 for numbers 1-11).

The actual bets offered may be different than this example, and you'll make several choices like this. Read the description of the 2 bets carefully each time before making your choices.

You would be doing us a favor if you chose the same bet in all 20 choices.

→

Powered by Qualtrics 



Caltech

Question 1 of 4:

- For each choice, the computer will draw a random number, 1-20.
- For each choice, you can choose either Bet A or Bet B. These are described below.

Bet A: You receive \$5 if the number drawn is from 1-16, and \$25 if it is from 17-20.	Bet B: You receive \$25 if the number drawn is from 1-16, and \$5 if it is from 17-20.																																								
<p style="text-align: center;">\$5</p> <table border="1"><tr><td>1</td><td>2</td><td>3</td><td>4</td></tr><tr><td>5</td><td>6</td><td>7</td><td>8</td></tr><tr><td>9</td><td>10</td><td>11</td><td>12</td></tr><tr><td>13</td><td>14</td><td>15</td><td>16</td></tr></table> <p style="text-align: center;">\$25</p> <table border="1"><tr><td>17</td><td>18</td><td>19</td><td>20</td></tr></table>	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	<p style="text-align: center;">\$25</p> <table border="1"><tr><td>1</td><td>2</td><td>3</td><td>4</td></tr><tr><td>5</td><td>6</td><td>7</td><td>8</td></tr><tr><td>9</td><td>10</td><td>11</td><td>12</td></tr><tr><td>13</td><td>14</td><td>15</td><td>16</td></tr></table> <p style="text-align: center;">\$5</p> <table border="1"><tr><td>17</td><td>18</td><td>19</td><td>20</td></tr></table>	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
1	2	3	4																																						
5	6	7	8																																						
9	10	11	12																																						
13	14	15	16																																						
17	18	19	20																																						
1	2	3	4																																						
5	6	7	8																																						
9	10	11	12																																						
13	14	15	16																																						
17	18	19	20																																						

Choice 1	Choice 2	Choice 3	Choice 4	Choice 5
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>
Choice 6	Choice 7	Choice 8	Choice 9	Choice 10
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>
Choice 11	Choice 12	Choice 13	Choice 14	Choice 15
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>
Choice 16	Choice 17	Choice 18	Choice 19	Choice 20
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>

If this screen is chosen for payment, the computer will draw 20 random numbers, one for each of

your 20 choices. After this, the computer will draw another number 1-20 to determine which of your 20 choices you'll actually be paid for.

I confirm these are the choices I want.



Powered by Qualtrics [↗](#)