CORRELATION NEGLECT IN STUDENT-TO-SCHOOL MATCHING

ALEX REES-JONES, RAN SHORRER, AND CHLOE TERGIMAN

ABSTRACT. A growing body of evidence suggests that decision-makers fail to account for correlation in signals that they receive. We study the relevance of this mistake in students' interactions with school-choice matching mechanisms. In a lab experiment presenting simple and incentivized school-choice scenarios, we find that subjects tend to follow optimal application strategies when schools' admissions decisions are determined independently. However, when schools rely on a common priority—inducing correlation in admissions—decision making suffers: application strategies become substantially more aggressive and fail to include attractive "safety" options. We document that this pattern holds even within-subject, with significant fractions of participants applying to different programs when correlation is varied but all payoff-relevant elements are held constant. We provide a battery of tests suggesting that this phenomenon is at least partially driven by correlation neglect, and we discuss implications that arise for the design and deployment of student-to-school matching mechanisms.

Keywords: Matching, correlation neglect, laboratory experiment.
JEL Codes: C91, D01, D03, M21.

Date: **January 2020**.

Rees-Jones: Cornell University, arr34@cornell.edu. Shorrer: Pennsylvania State University, shorrer@psu.edu. Tergiman: Pennsylvania State University, cjt16@psu.edu. For helpful comments, we thank Bnaya Dreyfuss, Rustam Hakimov, Ori Heffetz, Dorothea Kübler, Xiao Lin, Assaf Romm, Alex Teytelboym, and seminar participants at the ASSA Annual Meetings, New York University, and the Stanford Institute for Theoretical Economics. Shorrer was supported by a grant from the United States-Israel Binational Science Foundation (BSF grant 2016015).

A growing body of evidence suggests that many people struggle with decision-making in the presence of correlation. In typical examples of this problem, decision-makers are presented with multiple signals that are each influenced both by independent components and information from a common source. The process by which signals are generated induces correlation, and optimal decision-making requires taking it into account. In practice, however, experiments like those of Enke and Zimmermann (2019) demonstrate that many decision-makers neglect to do so, effectively acting as if these correlated signals are independent.

In this paper, we study the prevalence and consequences of these types of failures of reasoning in a decision of considerable importance: the application strategies of students applying to schools. Many application processes inherently require students to make forecasts of events determined by common underlying inputs, resulting in correlation structures like those described above. For example, students commonly must whittle a large number of schools down to a smaller set that are applied to or ranked, introducing an incentive to avoid listing two programs with highly correlated admissions decisions.[1] In such environments, a student harboring correlation neglect faces a challenging decision.

To illustrate these difficulties, consider a simple example. Imagine there are three programs at which you could potentially match, offering payoffs of 3, 2, and 1. Call these programs the best, middle, and worst programs, respectively. These programs will all rank you based on a common, currently-unknown, priority score; assume it will be a random integer drawn from a uniform distribution ranging from 0 to 99. The best program will admit you if your score is at least 50. The middle program will admit you if your score is at least 45. The worst program will admit you with any score. If you could only apply to two of these programs, to which two would you apply?

When considering this problem, one might feel the temptation to apply to the two programs with the highest payouts—we will refer to this as the *aggressive application strategy*. However, doing so is costly in expectation. Because these programs rely on the same score and have near-identical thresholds, the probability of being accepted by the middle program *conditional on being rejected by the best program* is quite low (10%), and insufficient to motivate a risk-neutral decision-maker from taking the sure payoff offered by applying to the worst school. Expected payout is maximized by applying to the best and the worst programs—we will refer to this as the *diversified application strategy*.

Consider next a slightly modified example. Imagine you are considering the same three programs, but now these programs rank you based on program-specific, statistically independent priority scores. Again, these evaluations are drawn from a uniform distribution ranging from 0 to 99. The best program's score threshold remains at 50, and the worst program continues to admit anyone. However, the middle program's acceptance threshold is changed to 90. In this situation, to which two schools would you apply?

[1]For a recent discussion of optimal diversification strategies in these environments—and their significant complexity—see Shorrer (2019).

As above, applying to the best and the worst programs remains the expected-value-maximizing strategy. Moreover, the consequences of pursuing either the diversified or aggressive application strategies are exactly the same as in the first example. The diversified application strategy grants a 50% chance of enrollment at the best program and a 50% chance of enrollment at the worst program. The aggressive application strategy grants a 50% chance of enrollment at the best program and, conditional on rejection there, a 10% chance of enrollment at the middle program. If one is restricted to these two strategies, choices across these scenarios should be identical.[2]

As we document in this paper, students' application strategies across these scenarios are quite different. When outcomes are correlated, a substantial fraction of students apply to the two most selective programs—i.e., they apply aggressively. By contrast, when priorities are determined independently, subjects' intelligently pursue the diversified application strategy at a much higher rate. Despite the numerical equivalency of probabilities, subjects act as if a 10% conditional probability of acceptance (the relevant probability for decision-making in the first example) is much more appealing than a transparent 10% unconditional probability of acceptance (the relevant probability for decision-making in the second example).

These simple examples illustrate something we believe to be a pervasive feature of school choice. In many environments, students can only apply to a subset of the schools that they see as attractive. In such situations, correlation in evaluations at different programs may be neglected or underweighted, leading students to fail to apply the appropriate contingent reasoning when deciding whether to apply to programs of similar selectivity.[3] The consequence is inadequate diversification of application portfolios conceptually similar to the inadequate diversification of asset holdings that is attributed to correlation neglect in Eyster and Weizsäcker (2016).

Concern about decision-quality in the face of correlated admissions is more than academic. As we summarize in Section 1, many countries use school assignment systems that involve choices much like the scenario just considered. In these systems, students are required to submit constrained lists of applications before discovering the results of the common test used to determine priority. To the extent that students are unable to correctly reason in such environments, intervention and revision of these systems may be merited.

To study this issue systematically, we ran a laboratory experiment among 165 students of Penn State University in early 2019. Subjects faced incentivized application scenarios much like the example above. In each scenario, subjects provided a rank-order list (ROL) to be used to match them to one of three schools. These ROLs could only contain two items, however, and thus required the student to choose a school application to forego.

To study the role of correlation neglect, we generate both within-subject and between-subjects variation in the correlation of admissions decisions. The presence or absence of correlation was

[2]Furthermore, while we have emphasized expected-value-maximizing behavior in our example, this equivalence is expected to hold more broadly. Indeed, it should hold so long as preferences depend only on the induced probability distribution over final matches.

[3]For detailed experimental examination of the difficulties of contingent reasoning in the presence of uncertainty, see Martínez-Marquina et al. (2017).

governed by whether programs' priorities were determined by either a single, common priority score or by program-specific, independent priorities, respectively. Subjects began the experiment by completing nine scenarios across which the acceptance thresholds varied but with the same correlation condition. They then completed a second battery of the nine scenarios presented under the other correlation condition. Comparing choices across these conditions provides a clean between-subjects analysis of how application strategies respond to correlated evaluation. A perhaps stronger test arises from the within-subject examination of "matched pairs" of scenarios, in which given strategies result in the same distribution of outcomes in both conditions.

Across these scenarios we document a systematic and quantitatively large tendency to fail to pursue the diversified strategy in the presence of correlation, in effect foregoing desirable "safety" options. We present two classes of evidence suggesting that this relates to incorrect processing of correlated environments. First, to provide a benchmark for correct processing, we present subjects with direct choices over monetary gambles. These monetary gambles were constructed to match the gamble induced by different application strategies in the scenarios seen by the subject. We find that choices in this transparent domain rationalize the choice of the diversified strategy, and are substantially more predictive of subjects' application strategies when they are made in the absence of correlation. Second, we present subjects with a variant of the experimental elicitation of correlation neglect of Enke and Zimmermann (2019). We find that this variable predicts subjects' propensity to switch between the diversified and aggressive application strategies in reaction to the correlation of admissions decisions.

This paper contributes to two literatures. First, and most directly, our paper contributes to the literature on correlation neglect. Common lab-experimental tests of correlation neglect (e.g., Enke and Zimmermann 2019) provide compelling evidence of the underlying behavioral bias. However, in order to isolate the role of correlation and in the interest of being maximally general, these tests are based on abstract forecasting tasks that are several steps removed from most field behaviors of interest. We contribute by identifying a way in which these abstract ideas become concretely relevant for a field behavior of substantial economic importance. We identify a class of large-scale matching systems of interest, provide theory tailored to understanding these environments, and provide tests that directly confirm the application is reasonable. We view this context as a conceptual proving ground for the field relevance of correlation neglect, and our experimental tests to confirm the need for the integration of these ideas into market design.[4]

Second, this paper contributes to a recently growing literature in "behavioral market design." While work in market design has typically assumed that market participants behave optimally, recent studies from both the lab[5] and the field[6] have shown a meaningful propensity of suboptimal behavior. While such studies suggest a role for behavioral economics in the modeling of matching

[4]Note that applications to school choice are not the only suggested field applications of correlation neglect; see also Eyster and Weizsäcker (2016) for applications to financial decision-making and Levy and Razin (2015) for applications to voting behavior.

[5]See Featherstone and Niederle (2016); Guillen and Hakimov (2017) and (2018); Ding and Schotter (2017); Basteck and Mantovani (2018); Li (2017); or Koutout et al. (2018).

[6]See Hassidim et al. (2018); Rees-Jones (2018); Rees-Jones and Skowronek (2018); and Shorrer and Sóvágó (2018).

markets, they provide relatively little guidance on the form such models should take. The examples cited above may be characterized as rejections of the null hypothesis of optimal reporting and reduced-form examinations of the correlates of mistakes; these studies do not isolate the fundamental biases that drive these decisions. This paper contributes by demonstrating the role of a specific behavioral model capable of making precise in- and out-of-sample predictions about biased respondents' reporting patterns. Such results are necessary to provide theorists with a means of acting on the observation that market participants struggle in these environments.[7]

The paper proceeds as follows. Section 1 presents summaries of the matching environments that motivate our study and guide our experimental design. Section 2 theoretically formalizes correlation neglect and its consequence of aggressive application strategies. Section 3 presents details of the design and deployment of our experiment, and Section 4 presents results. Section 5 concludes by discussing several means of combating these biases with market design.

1. Motivating Matching Environments

We begin by describing a set of existing matching systems that help motivate our interest in correlation neglect. While some degree of correlation in admissions decisions is ubiquitous in school choice environments, we focus on a class of systems where the correlation structure is particularly stark: systems in which application deadlines occur before students learn their performance on standardized tests that determine their priority. To the extent that uncertainty in admission is driven by uncertainty about test performance, this structure results in substantial correlation in admissions outcomes, and ultimately induces a decision problem quite similar to the example considered in the introduction.

Below, we summarize three national school-choice systems with these features, chosen both for their link to our experiment and for the presence of evidence of mistakes in applications strategies.

1.1. **The United Kingdom: The Universities and Colleges Admissions Service.**

The vast majority of college admissions in the United Kingdom are organized by the "Universities and Colleges Admissions Service" (UCAS).[8] When participating in the system, aspiring students may apply for up to five courses of study. These applications are due by mid-January, although some courses impose earlier deadlines.[9]

[7]For other examples of experiments testing the role of specific behavioral models as accounts of mistakes in matching markets, see Pan (2019) or Dargnies et al. (2019) examining self confidence or Dreyfuss et al. (2019) examining expectations-based reference dependence. Note that the models considered in these papers satisfy the property discussed in Footnote 2, implying that they predict no differences in behavior across our correlated and uncorrelated environment (conditional on choosing one of the focal strategies constructed to have equivalent payoffs across environments).

[8]In 2018, 695,565 applications to undergraduate-degree level courses were received, resulting in 533,360 matches. https://www.ucas.com/data-and-analysis/undergraduate-statistics-and-reports/ucas-undergraduate-end-cycle-reports/2018-end-cycle-report

[9]Oxford, Cambridge, and courses in medicine, dentistry, and veterinary science have application deadlines in mid-October.

At the time of application, test scores that are used for admissions decisions are not yet available for most of the applicant pool—specifically, A-level exams for those currently finishing their secondary education will typically only be available the following August. Consequently, this system is designed to permit educational institutions to make offers of admission *contingent on scoring above a specified threshold when these tests are taken*. Nearly all offers take this form.[10]

Due to this structure, students ultimately face a decision of which contingent contracts to pursue. By the end of March, students will have heard back from all of their five choices. At this time they must specify a "firm" choice and an "insurance" choice, and decline all other offers. This is effectively a commitment to attend the firm choice if the conditions of admission are met. If the firm choice's conditions are not met, the student is considered for admission at the insurance choice. If the insurance choice's conditions are then not met, the student is unmatched. While some procedures are in place to assist students who are ultimately unmatched after test scores are observed, students are strongly incentivized to be matched through the straightforward application of this process.[11]

As is readily apparent, students make decisions in this environment facing substantial uncertainty about how they will be evaluated, with this evaluation being correlated across schools. In the first stage, the student must whittle the universe of possible applications one could submit into a list of merely five, understanding that programs will have some degree of similarity in the manner in which they assess the student's extant profile. In the second stage, once offers are received, the student must whittle this set of offers into only two, typically with both of the offers conditioning on a common test score.

Several patterns in application data call the wisdom of students' applications into question. In a review of the system conducted in 2011, the UCAS found that

> "Many highly qualified applicants apply only to a narrow range of very selective [higher education institutions] which find it difficult to differentiate between these applications. This leads each year to a number of candidates with excellent grades failing to gain a place."

Furthermore, they found that 42% of applicants applying before test scores are available list an insurance choice with conditions for admission that are at least as stringent as those for the firm choice, in effect guaranteeing that the student remains unmatched if admission to the firm choice is not secured. In short, there is a general appreciation that a substantial fraction of students are pursuing unwise applications strategies (see UCAS (2011) for full documentation).

1.2. **Ghana: The Computerized School Selection and Placement System.** In Ghana, applications to "Senior High School"[12] are organized through the Computerized School Selection and

[10]In 2018, only 7.1% of offers were unconditional.

[11]Beyond creating worries about the consequences of correlation neglect, this structure facilitates regression-discontinuity analysis of the consequences of admission (see Broecke 2012).

[12]Following 6 years of primary school and 3 years of junior high school.

Placement System (CSSPS). This system, and problems that arise in students' use of it, is carefully examined in Ajayi (2013); we summarize key elements here.

Since 2005, senior high school admission has been conducted with a deferred acceptance algorithm (Gale and Shapley 1962). Through this system approximately 350,000 students are matched into 700 senior high schools every year. When participating in this match process, students submit rank-order lists of school/program-track pairs. Priorities in the schools are determined by the students' performance on the Basic Education Certificate Exam, which has not yet been taken at the moment of rank-order list submission. After performance on this exam is observed, the algorithm is applied and admissions are announced.

If students were able to list complete preference orderings, the well-known strategy-proofness property of deferred acceptance (Dubins and Freedman 1981, Roth 1982) would absolve students of the need to forecast their admission probabilities at different schools, and thus absolve them of their need to account for correlation in admissions decisions. However, the CSSPS imposes limits on the number of programs a student may rank. Upon initial implementation, students could only rank three choices; this was expanded to four in 2007 and six in 2008. This limit, which binds for the majority of students, introduces strong incentives to mitigate the risk of rejection from all listed programs, and the optimal strategy for choosing a portfolio of ranked schools depends crucially on the correlation structure of admission decisions (Shorrer 2019). As shown in Ajayi (2013), a substantial fraction of students submit rank-order lists with features that are ruled out by out optimal behavior. For example, 92% of students ranked schools in an order different than their selectivity, creating situations where rejection by the "back-up" option is assured conditional on rejection by the higher choice option. Furthermore, there is evidence that students coming from low-performing schools are more likely to follow these unwise application strategies, suggesting that differences in interactions with the matching mechanism help contribute to the less desirable admissions for students in this group.

1.3. **Kenya: Secondary School Admissions.** In Kenya, admissions to secondary school occur through a matching mechanism similar to those described above. At the end of eight years of primary-school education, students register and take the national Kenya Certificate of Primary Education (KCPE) examination. As part of the registration process—and crucially, prior to taking the exam—students submit their rank ordering over secondary schools. Government-run secondary schools are grouped into three quality-differentiated tiers: national, provincial, and district schools. Students list up to two choices from each tier, and are admitted to their most preferred option in the highest tier where they may be granted admission. Admissions depends on the outcome of the KCPE test as well as district quotas. For further discussion (on which this summary is based), see Lucas and Mbiti (2012).

As documented in Lucas and Mbiti (2012), patterns similar to those in the UK and Ghana arise. Among the top 5% of students in the 2004 administrative records of the KCPE—i.e., those with a realistic chance of admission to a national-tier school—36% of students listed a second choice school that was more selective than their first choice. As in the UK example, since the second

choice is only considered if admission to the first choice is denied, this pattern of reporting effectively foregoes one of only two opportunities for admission at this tier of school. Students making this error reduced their probability of admission to a national-tier school by over two percentage points—a large effect compared to the base admissions rate of 7.2% in the considered sample. While encouraging further study, Lucas and Mbiti conclude that "school choice errors in the admissions process could undermine offering the best opportunities to the highest ability students and cause inequalities to persist."

1.4. **Summary.** We have highlighted three large-scale matching systems with a key feature of interest: requirements to apply to a short list of schools in the presence of substantial uncertainty about a common factor affecting admissions. We note, however, that this structure is not limited to these three domains. Similar matching systems exist in China, Hungary, Trinidad and Tobago, and numerous sub-national applications. In short, this decision-environment is relatively common.

Despite the red flags raised about decision making in these systems, fully assessing the quality of application strategies is challenging. The evidence of mistakes summarized above is limited to cases where a subject lists an option with zero probability of realization—a mistake, to be sure, but a very extreme subclass of all mistakes. In field data, focus on mistakes like these is necessary because the analyst cannot observe the perceived utility associated with different schools. The absence of this data allows one to explain many questionable application strategies with extreme preferences rather than erroneous probability assessments. By contrast, when presenting subjects with experimental scenarios, the subjective value of different schools may be better controlled. When combined with the experimental manipulation of the degree of correlation in admissions decisions, this allows for the precise identification of the errors in reasoning we seek to study.

2. A Theory of School Applications with Correlation Neglect

In this section we formalize our discussion of correlation neglect. We state precisely its meaning in the context of school choice, then we establish its consequences for subjective expected utility and for preference submission. In the interest of proceeding to our experimental results quickly, we present our propositions with only brief intuitive explanation and relegate all formal proofs to Appendix A.

2.1. **Model preliminaries.** Consider a set X of schools. For simplicity, we assume that, conditional on the information available to students at the time of application, schools' admissions decisions are based on a single exam. Each agent has beliefs about how he will perform on the exam, summarized by the CDF F,[13] and he knows, for any $s \in X$, his utility from attending this school, u_s, and the score threshold required for admission, c_s. The utility from being unassigned is normalized to zero.

[13]Unless otherwise mentioned, we assume, without loss of generality, that students beliefs about scores are uniform on the unit interval (otherwise, apply the probability distribution transform to all scores, including thresholds).

In the leading example in the introduction, $X = \{\text{best, middle, worst}\}$; $c_{\text{best}} = 50$, $c_{\text{middle}} = 45$, and $c_{\text{worst}} = 0$; and $u_{\text{best}} = 3$, $u_{\text{middle}} = 2$, and $u_{\text{worst}} = 1$. In this example, the agent believes that scores are distributed uniformly over the integers between 0 and 99.

A *rank-order list* (ROL) is an ordered list of schools. Upon submission of an ROL, a student is admitted to the highest-ranked school at which admission is granted. Such ROLs are formally used in centralized matching markets applying, e.g., the deferred acceptance algorithm or its variants. Furthermore, Shorrer (2019) observes that ROLs may be considered to implicitly exist in decentralized school-choice markets, in which students attend the best school that accepts them.[14] Given an ROL, r, and an integer, i, we denote by r^i the i-th ranked school on that ROL. If an ROL ranks school j higher than school k, we denote this relationship by $s_j \succ s_k$.

2.2. **Correlation Neglect, Expected Utility of ROLs, and Chosen Application Strategies.** In this decision problem, we assume that students' evaluate the value of an ROL using standard subjective expected utility. Formally, this subjective expected utility is governed by the equation

$$\sum_{s \in r} Pr(\text{rejection at all } s_i \succ s) \cdot Pr(\text{admission at } s | \text{ rejection at all } s_i \succ s) \cdot u_s \tag{1}$$

We consider two types of agents, differing in their assessment of subjective expected utility. As a benchmark, we consider *sophisticated* agents. By assumption, these agents correctly evaluate all probabilities in Equation 1. We contrast this type against *(correlation) neglectful* agents, who replace the term $Pr(\text{admission at } s | \text{ rejection at all } s_i \succ s)$ with the term $Pr(\text{admission at } s)$. When we refer to the *subjective expected utility* of the neglectful type, which we denote by $V_n(r)$, we refer his expected utility based on his misguided probability assessments. When we refer to the *experienced utility* of either type (as well as the subjective expected utility of the sophisticated type), we refer to correctly evaluated expected utility, denoted $V_s(r)$.

Given these definitions, our use of the term "correlation neglect" may best be understood not as a reference to a fundamental, underlying bias, but as a reference to a reduced-form phenomenon: cases where correlation in outcomes necessitates Bayesian updating and the relevant updating is not pursued. It is worth noting that several different underlying mistakes in reasoning could generate this behavior. For example, an agent who completely understands the relevant correlation structure may fail to see any need for contingent reasoning. Alternatively, an agent who completely understands contingent reasoning may fail to see that correlation exists. While both examples result in neglecting the consequences of correlation, they derive from quite different underlying misunderstandings. While prior research has worked to disentangle the specific errors underlying these probabilistic judgments (see, e.g., Levin et al. 2016), in our environment the distinction does not result in differing predictions.

[14]This holds since students will only attend lower-ranked (less desirable) schools if they are rejected by all higher-ranked (more desirable) schools. Consequently, optimal ROLs can be calculated using similar dynamic programming as applies to the centralized case, and lower-ranked schools should be chosen conditional on the student being rejected by all higher ranked schools.

Illustrating these definitions in the context of our leading example, consider the ROL (best, middle)—the aggressive application strategy. The neglectful agent's subjective expected utility from this ROL is V_n (best, middle) $= .5 \times 3 + (1 - .5) \times .55 \times 2 = 2.05$. His expected experienced utility, which is equal to the sophisticated type's expected utility, is V_s (best, middle) $= .5 \times 3 + (1 - .5) \times .1 \times 2 = 1.6$. Because V_n (best, middle) $\geq$ V_s (best, middle), the neglectful agent perceives a higher expected utility than he would if he were sophisticated. This relative optimism is not a coincidence, as we illustrate in Proposition 1.

Proposition 1. *For any school choice environment and for any undominated rank-order list r,* $V_s(r) \leq V_n(r)$.

Put simply, because the neglectful agent fails to account for the "bad news" that a rejection conveys about the as-yet-unknown test score, he overestimates his chances of admissions after such a rejection occurs. This results in an overestimation of the utility that can be expected.

We now fully define behavior that we wish to characterize and study: ROL choice that maximizes the notion of subjective expected utility just established. Define a *school choice environment* to be the set of available schools and a vector of corresponding admission thresholds, formally denoted by $E = (X, c)$. Let $r(k, u, E)$ denote the optimal size-k ROL for agent with preferences u in environment E.[15] When E and u are clear, we often just write $r(k)$. Similarly, $r_n(k, u, E)$ and $r_n(k)$ denote the perceived optimal ROL of the neglectful agent.

To illustrate in our leading example, we calculate the subjective expected utility associated with the three undominated ROLs:

$$\begin{aligned} V_n\,(\text{best, middle}) &= .5 \times 3 + (1 - .5) \times .55 \times 2 && = 2.05 \\ V_n\,(\text{best, worst}) &= .5 \times 3 + (1 - .5) \times 1 \times 1 && = 2 \\ V_n\,(\text{middle, worst}) &= .55 \times 2 + (1 - .55) \times 1 \times 1 && = 1.55 \end{aligned}$$

Thus, the neglectful agent will chose (best, middle) over (best, worst) and (middle, worst). Formally, $r_n(2) =$ (best, middle). However, experienced utility is given by:

$$\begin{aligned} V_s\,(\text{best, middle}) &= .5 \times 3 + (1 - .5) \times .1 \times 2 && = 1.6 \\ V_s\,(\text{best, worst}) &= .5 \times 3 + (1 - .5) \times 1 \times 1 && = 2 \\ V_s\,(\text{middle, worst}) &= .55 \times 2 + (1 - .55) \times 1 \times 1 && = 1.55 \end{aligned}$$

.

[15] For simplicity, we assume that both the sophisticated and the neglectful type have a unique optimal ROL. This assumption, which is satisfied generically, plays no role in the analysis, and is only made simplifies statements.

Choosing an ROL to maximize V_n thus guides the agent to choose the aggressive application strategy (best, middle) when the diversified application strategy (best, worst) is objectively utility maximizing. The agent is expected to lose .4 experienced utils due to this mistake.

We next demonstrate that the consequences of correlation neglect for experienced utility may be grave. To do so, we first define a notion of the price of neglect that captures the fraction of experienced utility lost by the neglectful type.

Definition 1. *The price of neglect for a neglectful agent with utility u in environment E subject to constraint k is equal to the the difference in experienced utility between the maximizing size-k ROL and the ROL chosen by the neglectful type, normalized by the expected experienced utility from the maximizing ROL. In formal notation, $PN(u, E, k) = \frac{V_s(r(k)) - V_s(r_n(k))}{V_s(r(k))}$.*

Applying this definition, in the worst case, an optimal ROL of size k may generate k-times more experienced utility than the one that maximizes the subjective expected utility.

Proposition 2. *For any integer k, and any decision environment where the agent is constrained to (costlessly) apply to up-to-k schools, the price of neglect for the neglectful type is bounded above by $1 - \frac{1}{k}$. Furthermore, this bound is tight—for any k, there exist school choice environments where the price of neglect is arbitrarily close to $1 - \frac{1}{k}$.*

To illustrate how this worst-case bound may be achieved, consider a modification to our leading example. In this modification we add one more school to the choice set, and this school yields the same utility and has the same admissions cutoff as the best school. The neglectful agent would treat this copy of the best school as another (independent) chance for admission, ignoring the fact that rejection by one copy guarantees rejection by the other. As a result, the second application on his ROL is wasted, and he is no better off than he would be applying to a single school. As we show in Appendix A, for any permitted length of ROLs (k), we can construct examples involving perfect substitutes in which the neglectful agent will apply in a way that makes him no better off than if he had an ROL of length 1. Furthermore, a sophisticated agent can achieve approximately k times higher utility because the optimal length-k ROL in the examples we construct achieves approximately k times the utility of the optimal length-1 ROL.

While these extreme examples rely on the existence of perfect substitutes, note that the common situation of imperfect substitutes generates a similar conceptual force. In our leading example, because the best and middle programs have very similar cutoffs for admission, a rational agent should be hesitant to apply to both, and the failure to see this reasoning drives the utility losses documented above.

Given these observations, we conclude with a final result formally establishing the sense in which neglectful application strategies are overly aggressive.

Proposition 3. *For any constraint on the size of the ROL k, the neglectful type is at least as likely to be unassigned as the sophisticated type.*

In our leading example, recall that the sophisticated type would submit the ROL (best, worst), whereas the neglectful type would submit the ROL (best, middle). Because the worst school guarantees admission in the example, the sophisticated type faces no risk of being unassigned. By contrast, the neglectful type faces a 45% chance of remaining unassigned. In our example, this is characterized as an aggressive strategy: it contains options with higher utility, conditional on assignment, at the cost of exposing the agent to a greater degree of downside-risk of remaining unassigned. Proposition 3 demonstrates that the pursuit of more aggressive strategies is not unique to the example, but rather a general feature of ROL choice among neglectful agents.

2.3. **Summary.** The presence of correlation neglect leads subjects to be overly optimistic about their chances of admission to schools that they rank below their first choice. As a result, they undervalue the need to diversify the portfolio of schools contained in their ROL, resulting in overly aggressive application strategies. In the sections that follow, we test this prediction in the context of a laboratory experiment.

3. Experimental Design and Deployment

3.1. **Summary of Experimental Materials.** In this section, we summarize all measures and manipulations included in our lab experiment. All experimental materials are available in the Materials Appendix, and the structure of the experiment is summarized in Figure 1.

The experiment began with a brief informed consent document. Subjects were then told that the experiment was divided into parts (which we will refer to as modules), and that decisions in any part would not affect the opportunities presented in any other. Throughout the experiment, paper instructions were distributed and read out loud by the experimenter and subjects were given the opportunity to ask clarifying questions. The relevant experimental elicitations were then presented through a Qualtrics interface.

3.1.1. *Incentivized School-Choice Scenarios.* After reviewing the introductory materials, subjects were presented with the school-choice scenarios of primary interest. Within each scenario, students faced three programs to which they could apply. We referred to these programs as Colleges A, B, and C, and subjects could match to no more than one of them. To dictate the desirability of matching to these programs, each yielded a different payoff to matriculating students. Students matriculating to A, B, and C would receive a bonus payment of \$10, \$5, and \$2.5, respectively. Assignment to programs was determined by a matching procedure that depended on both test scores and students' rank-order lists (ROLs).

Test scores were simulated with draws from a uniform distribution ranging from 0 to 99, a structure known to participants. In the correlated-admissions module, a single test score is used

for all programs' admissions decisions. In the uncorrelated-admissions module, three statistically independent tests govern admissions to the three programs. In all cases, test scores are realized after the admissions lists are determined. Minimum test score requirements are presented as in Figure 2, and are varied across scenarios.

Based on this information, students were faced with the task of choosing a rank-order list (ROL) to be used in the assignment procedure. ROLs were ordered lists of two of the three schools; building an ROL required choosing one school application to forego and choosing an ordering among the remaining two applications that are submitted. The student was paired to the highest-ranked program for which the admission threshold was passed.

Construction of Scenarios: Our scenarios were constructed to function as "matched pairs," under which equivalent payoff structures were induced either in a correlated or uncorrelated decision environment. Table 1 summarizes each matched pair of scenarios and presents the gamble induced by the two focal ROLs. To illustrate, consider the first scenario. When outcomes were determined by a single priority (i.e., in the correlated-admissions module), the first scenario had a score threshold of 50 for school A, 45 for school B, and 0 for school C. When outcomes were determined by multiple, independent priorities (i.e., in the uncorrelated-admissions module), the thresholds were 50, 90, and 0. This example matches the scenario presented in the opening paragraphs of the paper. The aggressive application strategy ($(A \succ B)$) and the diversified application strategy ($(A \succ C)$) result in the same probabilities of admissions at each school and thus equivalent expected payouts, summarized in right columns of the table.

We constructed our nine matched pairs with several considerations in mind. While we were initially motivated by pilot results arising from scenario 1, we wanted to ensure that the patterns of behavior we observed were not somehow unique to the thresholds in that scenario.

First, we were concerned that some of the applications to program C might be due to the attraction of a completely certain option. This motivated the creation of scenario 2, which closely mirrors scenario 1 but makes admissions to the worst program uncertain (but still very likely). As we vary other thresholds across additional scenarios, we continue to create pairs that differ only in the certainty of admission to the third program (see scenario pairs (3, 4), (6,7) and (8,9)).

In scenarios 1 and 2 (as well as all other scenarios we will discuss), pursuing the aggressive application strategy ($(A \succ B)$) induces a gamble that is both riskier and (weakly) lower in expected value than the diversified application strategy ($(A \succ C)$). In scenarios 3 and 4, we set the test score thresholds in order for the aggressive application strategy to yield a higher expected value, making it potentially desirable for some risk preferences.

We constructed scenario 5 to study the extreme type of mistakes observed in the field settings described in Section 1: submitting second-choice options for which rejection is guaranteed conditional on rejection by the first choice. In the correlated-admissions module, the required test score for the middle program was 55, whereas the required score for the best program was 50. In the uncorrelated-admissions module, the independent priority score necessary for admission to the second program was 100. In both modules, applying to the top two programs yields a 50%

chance of admission to the best program and a 0% chance of admission to the middle program. Submission of the preference order $(A \succ B)$ therefore mirrors the worrying behaviors seen in the UK, Ghana, and Kenya.

Note that all of scenarios 1-5 are designed to focus on the pursuit of ROLs $(A \succ B)$ and $(A \succ C)$: the remaining undominated preference order $((B \succ C))$ is not meant to be appealing and empirically is rarely chosen. Scenarios 6-9 were included to examine application behavior in cases where the ROL $(B \succ C)$ is made more attractive (although our focus remains on preferences $(A \succ B)$ and $(A \succ C)$).

Across these scenarios, we may examine how the pursuit of the aggressive and diversified application strategies responds to correlation when a battery of other considerations are varied. These scenarios were divided into two blocks of nine, with the correlation structure constant within each block and the order of questions randomized within block at the subject level.

3.1.2. *Auxiliary Measures and Questions.* Following the school choice scenarios, three additional groups of questions were presented.

Preferences over Lotteries: First, subjects were presented with a series of nine choices over risky gambles. These gambles were constructed to match the gamble over monetary outcomes induced by the two focal admissions strategies ($(A \succ B)$ and $(A \succ C)$) submitted in each of the nine school-choice scenarios, as seen in Table 1. By eliciting direct preferences over these gambles, we may benchmark the choices made in the school-choice scenarios against choices that are made when their consequences are fully transparent.

Raven's Matrices: Second, subjects were presented with a battery of "Raven's Progressive Matrices," a common assessment of spatial reasoning used as a proxy for general cognitive ability (Raven and Raven 2003). Subjects were given five minutes to complete as many of the 6 matrices as they could.

Direct Elicitation of Correlation Neglect: Third, subjects faced a correlation-neglect elicitation closely derived from that used by Enke and Zimmermann (2019). Subjects were given the task of forecasting an underlying value, denoted "X." X is drawn from a normal distribution with a mean of zero and a standard deviation of 500. Subjects are asked to guess the value of X and are compensated based on their accuracy. The probability of winning a $10 bonus is governed by the squared difference between the subjects' estimate and the true value, providing incentives for truthful reporting. To help guide this decision, four noisy signals of the true value are drawn, and are communicated to the subject by "communication machines" (CMs). The functioning of these CMs induces correlation into their communicated forecast: one of the four signals is observed by all four CMs. One CM directly reports the common signal, whereas the other three report the average of the common signal and a signal only observed by that CM. All details of the noise distributions, signal generation, and reporting structure are communicated to participants.

As shown by Enke and Zimmermann (2019), this environment offers a clear way to measure the degree of correlation neglect. The optimal forecast in this environment is constructed by using the known correlation structure to back out the 4 signals provided to the CMs, then average those

signals. Denote this optimal forecast as f^o. Alternatively, one could imagine a subject treating the four reports of the CMs as if they were four independent signals and simply averaging them. Denote this naive forecast as f^n. As long as $f^n \neq f^o$, any individual forecast maps onto a specific value of χ implicitly defined by the equation $f = \chi f^n + (1-\chi)f^o$. Enke and Zimmermann use χ as a measure of the degree of correlation neglect in subjects' forecast, noting that $\chi = 0$ corresponds to a completely optimal forecast and $\chi = 1$ corresponds to treating the data as if it were independent.

To help mitigate measurement error, we follow Enke and Zimmermann's strategy of offering subjects multiple forecasting tasks and assigning them the median of their measured χ values. Subjects completed 3 forecasts.

3.1.3. *End of experiment.* The experiment concluded with a brief elicitation of demographics. Following these questions, bonuses for incentivized modules were determined and final payments for the experiment were made.

3.1.4. *Compensation.* Subjects received a show-up fee of \$7.00. In addition, subjects were incentivized to truthfully report their preferences and carefully answer questions. The incentives for each module were explained prior to its presentation, and are summarized here:

- One round was randomly chosen among the school choice scenarios and the equivalent preferences over lotteries questions. If that randomly chosen round was one of the school choice scenarios, the student's score(s) was compared with the submitted ROL. Earnings consisted of the bonus associated with the school that a student was admitted to. If the randomly chosen round was one of the lottery questions, we ran the lottery that the subject selected, and earnings consisted of the outcome of the selected lottery.
- Subjects also received \$1.00 for each correctly solved Raven's Matrix.
- One of the three direct correlation neglect elicitation questions was randomly selected. A subject received an additional \$10 if the submitted forecast was "close enough" to the true value.[16]

Subjects were informed of their earnings in each module only at the end of the session.

Average total earnings were \$18.10, and ranged from \$7.00 to \$33.00 (including the \$7.00 show-up fee).

3.2. **Preregistration.** We preregistered our hypotheses of interest, primary analyses, target sample size, and inclusion rules prior to the beginning of data collection. Preregistration documents are archived on aspredicted.org, and are included in the Materials Appendix for ease of reference. As we report results, we flag any and all places where we deviate from our preregistered analysis plan.

[16]As in Enke and Zimmermann (2019), the threshold for "close enough" was determined by a random draw.

3.3. **Experimental Deployment.** Our experiment was conducted in January and February 2019 at the Laboratory for Economics, Management and Auctions (LEMA) at Penn State University. Basic demographics are presented in Table 2. Subjects participated in one session only. We recruited 80 subjects in the "correlated admissions first" treatment and 85 in the "uncorrelated admissions first" treatment, consistent with our preregistered target of 80 per cell. The treatments differed only in whether students saw the correlated or uncorrelated questions first.[17]

4. Experimental Results

4.1. **Application Strategies in Scenarios.** Table 3 summarizes the application strategies pursued in each of our scenarios. Examining the first row, we see that when the scenario considered in the introduction is presented with correlated admissions decisions, 44.9% of subjects pursued the diversified application strategy ($(A \succ C)$). Among students submitting another ROL, by far the most common was $(A \succ B)$—the aggressive application strategy. 48.5% of students pursued this strategy despite its greater risk and lower expected value.[18] This behavior—which we interpret as a mistake—is substantially less prevalent in the uncorrelated-admissions module. In this environment, only 10.9% of students submitted $(A \succ B)$, with 84.2% of students making the optimal choice of $(A \succ C)$. In short, in this scenario, subjects were more tempted to pursue the aggressive (and perhaps unwise) admissions strategy of $(A \succ B)$ in the correlated admissions environment.

To assist in assessing these claims statistically, the final three columns of the table present p-values arising from a set of cross-module hypotheses tests. The first column presents a Fisher's exact test of whether the distribution of chosen ROLs varies by correlation module. The second and third columns present two-sample difference-of-proportions tests of equality in the fraction choosing the aggressive and diversified application strategies. Examining these statistics for the first scenario demonstrates that all differences discussed in the prior paragraph are unlikely to arise by chance.

Across the different thresholds induced across the nine scenarios, the patterns described above always holds: the strategy of diversifying to the best and the worst programs is more likely to be pursued in the uncorrelated-admissions module, and our target "tempting" behavior of aggressively applying to the top two programs is more likely in the presence of correlated evaluations. This remains true when the worst program still has residual uncertainty of admission (as in scenarios 2, 4, 7, and 9); when reducing risk by applying to the worst program comes at some cost in expected value (as in scenarios 3 and 4); when admission to the middle program is impossible conditional on rejection by the best program (as in scenario 5), and in environments in which the ROL $(B \succ C)$ is made more attractive (as in scenarios 6-9).

[17]Because we generally do not detect order effects (under which the distribution of, e.g., correlated choice would depend on whether the module appeared first or second), we present analyses which pool questions of the same type regardless of position in the experiment. In Appendix B we present analyses supporting this decision, including formal tests for order effects and a recreation of our main analysis using only data from the first module seen.

[18]Formally, the distribution of outcomes resulting from the ROL $(A \succ B)$ is second-order stochastically dominated by that resulting from $(A \succ C)$, meaning it should not be chosen by any risk-averse expected-utility maximizer.

Across these scenarios, we see a variety of patterns indicating responsivity to the incentives introduced by these variants. More subjects submit preference order ($A \succ B$) when it yields a higher expected value than the alternative (as in scenarios 3 and 4), fewer subjects submit ($A \succ B$) when it is a dominated strategy (as in scenario 5), and more subjects submit $(B \succ C)$ when it is made attractive (as in scenarios 6-9). And yet, across all these variants and despite these signs of intelligent response to incentives, substantially different patterns of reporting are seen based on the presence or absence of correlation in evaluation.

Figure 3 summarizes these differences by presenting the distribution of chosen ROLs averaged over all nine scenarios. On average, the rate of pursuit of the aggressive application strategy increases by 20.1 percentage points when scenarios are presented with correlated admissions decisions.

Patterns like these are additionally seen in within-subject evaluations. For each question, Table 4 shows the pairs of strategies submitted by each subject, focusing specifically on the two focal strategies of ($A \succ B$) and ($A \succ C$). The first two columns present the fraction of subjects choosing a focal strategy and not responding to the presence of correlation: these subjects pursue the same admissions strategy in both of the matched decisions. The next two columns present subjects who switched between focal strategies depending on the correlation condition. As illustrated in these columns, substantial fractions of subjects had their answer vary depending on the presence of correlation, with a strong tendency to chose the aggressive strategy ($(A \succ B)$) in the correlated-admissions module and the diversified strategy ($(A \succ C)$) in the uncorrelated-admissions module. Subjects switched from the aggressive strategy in the correlated-admissions module to the diversified strategy in the uncorrelated-admissions module in 21.1% of subject-scenario pairs, indicating a substantial propensity of this specific preference reversal. As is indicated in the final column, across most scenarios the majority of subjects who submitted the ROL ($A \succ B$) under correlation would submit the ROL ($A \succ C$) under independence. In contrast, the mirroring inconsistency (choosing ($A \succ C$) under correlation and ($A \succ B$) under independence) never exceeds 6.0%.

4.2. **Lottery choices.** In the experiment, subjects faced a series of questions in which they directly chose between the pairs of lotteries presented in the right-most columns of Table 3. These decisions directly elicit preferences over the transparent monetary consequences of submitting ($A \succ B$) or ($A \succ C$).

As shown in Figure 4, application strategies were substantially less aligned with gamble choices in the presence of correlation. In the correlated-admissions module, application decisions were consistent with the chosen transparent gamble in 44.0% of cases. In the uncorrelated-admissions module the rate of consistency between application strategies and chosen gambles rises to 66.9%, a nearly 23 percentage point increase.

Table 5 presents several additional statistics characterizing subjects' preferences over gambles. As illustrated in the first column, when payoffs induced by the aggressive and diversified strategy were made transparent, subjects overwhelmingly chose in favor of the latter. This supports our interpretation of that ROL as the optimal strategy with pursuit of other ROLs being misguided.

In the second column, we use this measure to help resolve which ROL should be honored among the subjects whose strategy differed across modules. Among subjects who chose ($A \succ B$) in the correlated-admissions module and ($A \succ C$) in the uncorrelated-admissions module—i.e., the individuals of interest in the previous subsection—we find that relatively few chose the gamble associated with ($A \succ B$). Only in scenario 4—one of the two in which the expected value of ($A \succ B$) exceeds that of ($A \succ C$)—do preferences on average tilt towards the aggressive strategy. In the third column in the table, we look at the ($A \succ C$) choices in the lottery questions for those subjects who consistently chose ($A \succ C$) in the school choice setting, providing validation from a third source of the consistency in their responses.

In summary, subjects' preferences over ROLs in the uncorrelated-admissions module tend to align with their preferences inferred from transparent presentations of the induced lotteries. This corroborates our treatment of the more aggressive application strategies in the correlated decision environment as arising from a mistake.

4.3. **Predicting the pursuit of aggressive strategies with the Enke and Zimmermann Correlation Neglect Measure.** To further validate the relationship between our behavior of interest and correlation neglect, we asked subjects 3 questions taken directly from the materials of Enke and Zimmermann (henceforth EZ; 2019). Following their technique, we directly compute their parameter of correlation neglect for each of those questions, and assign to each subject their median value. Because their parameter is meant to be interpreted on the unit interval—with a value of zero implying completely correct processing and a value of 1 implying that correlated signals are treated as perfectly independent—we restrict attention to cases where this measure falls in this range.[19]

Our goal in these analyses is to use the EZ measure to predict the propensity of the within-subject preference reversals that we argue are generated by correlation neglect. To that end, we construct a measure that compares the rate of optimal decisions versus the rate of the specific within-subject mistake that we considered in the prior section. Other anomalous ROL reporting patterns are effectively disregarded.[20]

Figure 5 presents a local-polynomial regression of the relative rate of our preference reversals of interest on the Enke-Zimmermann correlation neglect measure. As illustrated in this figure,

[19]We note that we did not preregister our intention to eliminate outlier values of the EZ measure because we did not anticipate their prevalence. In our data, 9% of subjects have an EZ measure below 0, ranging to a minimum value of -0.845. More worryingly, 31% of subjects have an EZ measures exceeding 1, ranging to a maximum value of 1.837. The frequency of these difficult-to-interpret values posed challenges for interpretation of our preregistered analyses, leading us to make these modifications to our plan.

[20]Formally, our dependent variable takes the value of 1 for a matched pair when the subject chose ($A \succ B$) in the correlated-admissions module and ($A \succ C$) in the uncorrelated-admissions module—the behavior that we attribute to correlation neglect—and zero when they chose ($A \succ C$) in both modules—optimal behavior. All other paths are ignored. Within-subject, we then calculate the fraction of cases in which the subject pursued the correlation-neglectful path. The mean (median) number of observations per subject used to calculate this fraction is 4.9 (5). Note that this measure is undefined for the 8 subjects who followed neither the optimal nor the correlation neglectful reporting pattern for any of the scenarios considered. Their exclusion leads to a usable sample of 157 observations for these analyses.

increasing from zero EZ correlation neglect to full EZ correlation neglect is associated with a substantial rise in the rate of choosing the aggressive strategy under correlation and the diversified strategy under independence (relative to the rate of consistently choosing the optimal strategy across correlation conditions).

To assess the statistical significance of this relationship and its robustness to controls, Table 6 presents analogous regression results. In the first column, we see that the EZ measure has a large, positive, statistically significant relationship with our measure of the propensity of preference reversals. This relationship persists with the inclusion of our Raven's Matrices measure of cognitive ability as well as the full battery of demographic variables collected in the study. Beyond the EZ measure, the only additional variable found to be predictive is cognitive ability (a variable commonly found to be associated with making mistakes in matching mechanisms, see Basteck and Mantovani (2018); Hassidim et al. (2018); Rees-Jones (2018); Rees-Jones and Skowronek (2018)).[21]

In summary, the key behavior we have attributed to correlation neglect—pursuing the aggressive application strategy when admissions decisions are correlated and the diversified strategy when they are not—is predicted by existing experimental measures of correlation neglect.

4.4. **Alternative Explanations.** As described in Section 3.1.1, we purposefully constructed the nine scenarios in our study to address a range of potential confounds. Our results in Section 4.1 generally provided little evidence of a role for these worries. In this section, we address several remaining commonly considered alternative explanations for our results.

4.4.1. *Aversion to Schools Dominated as Singleton Applications.* Recall that in our leading example (matching Scenario 1), subjects were more likely to submit an ROL rationally foregoing the middle school when that school had an admissions threshold of 90 on an independent test as opposed to when it had an admissions threshold of 45 on the common test. While both framings result in the same distributions of outcomes when the ROL (best, middle) are applied, note that they would result in different outcomes if middle were listed in isolation. In the uncorrelated framing, the middle option is formally dominated as a singleton application: it has a lower utility and a higher admissions threshold than the best school. In the correlated framing, it is not dominated: while it does have a lower utility, it also has a lower threshold. If subjects are irrationally averse to including such options in a multi-school ROL, this aversion could guide them towards optimal behavior (for the wrong reasons).

Note, however, that while this concern is present in our leading example, it is not present in scenarios 3, 4, 8, or 9. In all such cases, compared to applications to top school A, applications to middle school B yields a lower utility with a higher chance of admissions regardless of framing. The continued presence of qualitatively large differences in application behavior in these environments alleviates the worry that this potential aversion explains the results we have documented.

[21]We additionally preregistered that we would examine the relationship between our within-subject-mistakes DV and Raven's task performance without additional controls. In this analysis we find a similar negative relationship ($\beta = -0.059$, s.e.$= 0.024$, $p = 0.015$).

4.4.2. *Models of Choice-Set Dependence.* We consider next the potential explanatory value of a class of choice-set-dependent models of recent prominence in the behavioral economics literature. These models consider how the distribution of attributes in a choice set influence the decision weights placed on those attributes, with greater weight meant to capture the devotion of additional attention. In the focusing model of Köszegi and Szeidl (2012), it is assumed that an attribute with a larger range of values receives more decision weight. In the relative thinking model of Bushong et al. (2019), it is assumed that an attribute with a more narrow range gets more attention. In the salience model of Bordalo et al. (2012), the key predictions come from their assumptions of ordering and diminishing sensitivity, which at times point in the direction of either of the previous models. For a recent paper carefully comparing these theories and their empirical performance in explaining experimental purchasing decisions, see Somerville (2019).

When applying these frameworks to our setting, we believe it is most natural to imagine the student to be considering two attributes: payoff and admissions threshold. Payoffs are held constant in our design, but admissions thresholds (and their ranges) differ. Referencing Table 1, note that the manner in which thresholds change makes the range of thresholds in the uncorrelated treatment larger in some scenarios (1, 2, 5-7), smaller in other scenarios (4, 9), and unchanged in yet others (3, 8). The fact that we document apparent neglect of the safety option that is most attractive on the admissions-threshold dimension across all of these variants suggests that choice-set-dependent models based on comparisons of range do not provide a natural explanation for our results.[22]

4.4.3. *Preferences for Randomization.* We interpret our finding of within-subject preference reversals to be strong evidence of incorrect processing of correlated environments. This interpretation relies on the assumption that, were all elements of the decision environment fully understood, behavior would not respond to framing. However, several recent works have examined cases where subjects hold an explicit preference for randomization (see, e.g., Agranov and Ortoleva 2017, Dwenger et al. 2018, Cerreia-Vioglio et al. 2019); in the presence of such preferences, inconsistent choice need not reflect a true preference reversal.

Four pieces of evidence suggest that preference for randomization have little role in our results. First, note that when subjects faced their first module of school-choice decisions, they did not know that an additional round of equivalent-but-differently-framed scenarios would follow. Without knowledge that two iterations of each question would occur, the underlying motivations that guide intentional randomization would not be triggered. Second, we note that intentional randomization alone would not generate the stark asymmetry observed: while it could predict different answers within-subject, it would not predict the strong tendency for aggressive applications specifically under correlated framing. Third, a preference for randomization would not explain why choices made in the absence of correlation were more in line with choices in transparant gambles. Finally,

[22]Note that only the ordering assumption in Bordalo et al. (2012) depends directly on range. Turning to the assumption of diminishing sensitivity, we note that this assumption considers the salience implications of shifting the values of attributes to be larger across both options considered. Since most of our scenarios vary only the threshold for middle program while holding the thresholds for the other programs constant, this assumption does not apply to our setting as written.

even if a preference for randomization obfuscates the interpretation of within-subject preference reversals, the between subject contrasts we have presented would remain valid. Given these issues, we believe that preference for randomization does not provide a systematic account of the results we have documented.

5. Discussion

In this paper, we have noted that correlation neglect offers a natural explanation for some of the difficulties observed in a common class of school-choice problems: those in which rankings are submitted prior to essential test scores being determined. We have followed in the tradition of work such as Chen and Sönmez (2006) in using controlled lab experiments to directly assess students' response to school-assignment mechanisms. We note that our results are not unique to centralized markets that use such mechanisms, and pervade decentralized school admissions problems as well. Across the incentivized scenarios that we consider, we see clear evidence that experimentally manipulated correlation in admissions decisions results in overly aggressive application strategies, and that decisions improve when the environment is modified so that correlation is removed.

Difficulties in comprehension induced by correlation are broadly relevant for market designers. The systems we highlighted in Section 1 offer somewhat extreme examples of the forecasting challenges that we have described, but some degree of these challenges are ubiquitously present. Whether due to institutional constraints or scarce consideration time, students often have to choose a comparatively small set of schools to ultimately apply to or rank. Even in cases where all inputs to evaluation are known, students often harbor some uncertainty about how their profile will be evaluated. The fact that programs often have some agreement on the evaluation process results in correlation. In such environments, students' approach to the matching process might be meaningfully suboptimal not due to a failure to optimally rank-order their considered schools (the typical mistake of interest in behavioral matching papers, see Hakimov and Kübler (2019) for a recent review), but rather due to forming the wrong consideration set in the first place. And indeed, forming an unwise consideration set can be a substantially more costly mistake, since the primary risk induced is not matching to the wrong school (as arises from misordering preferences in deferred acceptance), but rather failing to match at all.

The presence of correlation neglect directly influences several practical considerations of market design. One immediate consequence is a connection between the literature on priority-tie-breaking in school choice to the discussion on protecting unsophisticated agents (e.g., Pathak and Sönmez 2008, Hassidim et al. 2017a, Rees-Jones 2017). Tie-breaking methods have been studied extensively, especially comparing the use of a single common lottery with multiple independent lotteries (e.g., Abdulkadiroğlu et al. 2009, Ashlagi et al. 2019).[23] Our results imply that, when applications are restricted or costly, there may exist a tension between efficiency (which may be improved by

[23]While this literature focused on centralized mechanisms, admission decisions are also independent in decentralized markets where oversubscribed schools use independent (school-specific) lotteries to break priority ties between applicants (e.g., Dobbie and Fryer 2011). Similarly, if all schools use a single lottery or exam to break priority ties, admission decisions are correlated.

the use of a single lottery) and the desire to protect unsophisticated agents. When concerns about protecting the unsophisticated become central, the use of multiple tie-breaking rules is recommended.

In cases where a market designer is soliciting a highly constrained rank-order list, our results suggest that it might be worthwhile to assist subjects with the contingent reasoning that correlated admissions necessitate. For example, consider a system in which a student is submitting an aggressive rank-order list. To the extent that correlation neglect is present, the market designer should worry that subject considered only the unconditional probability of acceptance of each program rather than the relevant probability of acceptance conditional on rejection by all higher ranked programs. In environments where admissions probabilities are available to the market designer, it is relatively easy to communicate the crucial probabilities for decision-making in a simple pop-up window. A simple pop-up intervention like this has been used to combat deviations from truthful preference reporting in the Israeli Psychology Match (Hassidim et al. 2017b, 2018), and the success seen there has motivated similar interventions in genetic matching services.

In summary, and to conclude, the specific failures in reasoning induced by correlation neglect directly interplay with crucial technical aspects of market design. Based on the magnitude of effects observed in our study, combined with the substantial prevalence of dominated choices seen in the matching systems reviewed in section 1, we believe that integration and accommodation of correlation neglect into our frameworks for market design can be of substantial benefit to the millions of students who interact with these systems.

REFERENCES

ABDULKADIROĞLU, A., P. A. PATHAK, AND A. E. ROTH (2009): "Strategy-Proofness versus Efficiency in Matching with Indifferences: Redesigning the NYC High School Match," American Economic Review, 99, 1954–78.

AGRANOV, M. AND P. ORTOLEVA (2017): "Stochastic Choice and Preferences for Randomization," Journal of Political Economy, 125, 40–68.

AJAYI, K. F. (2013): "School Choice and Educational Mobility: Lessons from Secondary School Applications in Ghana," Mimeo.

ASHLAGI, I., A. NIKZAD, AND A. ROMM (2019): "Assigning More Students to their Top Choices: A Comparison of Tie-Breaking Rules," Games and Economic Behavior, 115, 167 – 187.

BASTECK, C. AND M. MANTOVANI (2018): "Cognitive Ability and Games of School Choice," Games and Economic Behavior, 109, 156 – 183.

BORDALO, P., N. GENNAIOLI, AND A. SHLEIFER (2012): "Salience Theory of Choice Under Risk," Quarterly Journal of Economics, 127, 1243–1285.

BROECKE, S. (2012): "University Selectivity and Earnings: Evidence from UK Data on Applications and Admissions to University," Economics of Education Review, 31, 96–107.

BUSHONG, B., M. RABIN, AND J. SCHWARTZSTEIN (2019): "A Model of Relative Thinking," Mimeo.

CERREIA-VIOGLIO, S., D. DILLENBERGER, P. ORTOLEVA, AND G. RIELLA (2019): "Deliberately Stochastic," American Economic Review, 109, 2425–45.

CHADE, H. AND L. SMITH (2006): "Simultaneous Search," Econometrica, 74, 1293–1307.

CHEN, Y. AND T. SÖNMEZ (2006): "School choice: an experimental study," Journal of Economic Theory, 127, 202–231.

DARGNIES, M.-P., R. HAKIMOV, AND D. KÜBLER (2019): "Self-Confidence and Unraveling in Matching Markets," Management Science, 65, 5603–5618.

DING, T. AND A. SCHOTTER (2017): "Matching and chatting: An experimental study of the impact of network communication on school-matching mechanisms," Games and Economic Behavior, 103, 94–115.

DOBBIE, W. AND J. FRYER, ROLAND G. (2011): "Are High-Quality Schools Enough to Increase Achievement among the Poor? Evidence from the Harlem Children's Zone," American Economic Journal: Applied Economics, 3, 158–87.

DREYFUSS, B., O. HEFFETZ, AND M. RABIN (2019): "Expectations-Based Loss Aversion May Help Explain Seemingly Dominated Choices in Strategy-Proof Mechanisms," SSRN WP 3381244.

DUBINS, L. E. AND D. A. FREEDMAN (1981): "Machiavelli and the Gale-Shapley Algorithm," The American Mathematical Monthly, 88, 485–494.

DWENGER, N., D. KÜBLER, AND G. WEIZSÄCKER (2018): "Flipping a coin: Evidence from university applications," Journal of Public Economics, 167, 240–250.

ENKE, B. AND F. ZIMMERMANN (2019): "Correlation Neglect in Belief Formation," The Review of Economic Studies, 86, 313–332.

EYSTER, E. AND G. WEIZSÄCKER (2016): "Correlation Neglect in Portfolio Choice: Lab Evidence," SSRN WP 2914526.

FEATHERSTONE, C. R. AND M. NIEDERLE (2016): "Boston versus Deferred Acceptance in an Interim Setting: An Experimental Investigation," Games and Economic Behavior, 100, 353 – 375.

GALE, D. AND L. SHAPLEY (1962): "College Admissions and Stability of Marriage," American Mathematical Monthly, 69, 9–15.

GUILLEN, P. AND R. HAKIMOV (2017): "Not quite the best response: truth-telling, strategy-proof matching, and the manipulation of others," Experimental Economics, 20, 670–686.

——— (2018): "The effectiveness of top-down advice in strategy-proof mechanisms: A field experiment," European Economic Review, 101, 505–511.

HAKIMOV, R. AND D. KÜBLER (2019): "Experiments on matching markets: A survey," Tech. rep.

HASSIDIM, A., D. MARCIANO, A. ROMM, AND R. I. SHORRER (2017a): "The Mechanism Is Truthful, Why Aren't You?" American Economic Review, 107, 220–24.

HASSIDIM, A., A. ROMM, AND R. I. SHORRER (2017b): "Redesigning the Israeli Psychology Master's Match," American Economic Review, 107, 205–09.

——— (2018): "'Strategic' Behavior in a Strategy-Proof Environment," SSRN WP 2784659.

KÖSZEGI, B. AND A. SZEIDL (2012): "A Model of Focusing in Economic Choice," The Quarterly Journal of Economics, 128, 53–104.

KOUTOUT, K., A. DUSTAN, M. VAN DER LINDEN, AND M. WOODERS (2018): "Mechanism Design under Strategy Advice and Sub-Optimal Play: A School Choice Experiment," Mimeo.

LEVIN, D., J. PECK, AND A. IVANOV (2016): "Separating Bayesian Updating from Non-Probabilistic Reasoning: An Experimental Investigation," American Economic Journal: Microeconomics, 8, 39–60.

LEVY, G. AND R. RAZIN (2015): "Correlation Neglect, Voting Behavior, and Information Aggregation," American Economic Review, 105, 1634–45.

LI, S. (2017): "Obviously Strategy-Proof Mechanisms," American Economic Review, 107, 3257–87.

LUCAS, A. M. AND I. M. MBITI (2012): "The Determinants and Consequences of School Choice Errors in Kenya," American Economic Review, 102, 283–88.

MARTÍNEZ-MARQUINA, A., M. NIEDERLE, AND E. VESPA (2017): "Probabilistic States versus Multiple Certainties: The Obstacle of Uncertainty in Contingent Reasoning," Working Paper 24030, National Bureau of Economic Research.

PAN, S. (2019): "The instability of matching with overconfident agents," Games and Economic Behavior, 113, 396 – 415.

PATHAK, P. A. AND T. SÖNMEZ (2008): "Leveling the Playing Field: Sincere and Sophisticated Players in the Boston Mechanism," American Economic Review, 98, 1636–52.

RAVEN, J. AND J. RAVEN (2003): Raven Progressive Matrices, Boston, MA: Springer US, 223–237.

REES-JONES, A. (2017): "Mistaken Play in the Deferred Acceptance Algorithm: Implications for Positive Assortative Matching," American Economic Review, 107, 225–29.

——— (2018): "Suboptimal Behavior in strategy-Proof Mechanisms: Evidence from the Residency Match," Games and Economic Behavior, 108, 317 – 330.

REES-JONES, A. AND S. SKOWRONEK (2018): "An Experimental Investigation of Preference Misrepresentation in the Residency Match," Proceedings of the National Academy of Sciences, 115, 11471–11476.

ROTH, A. E. (1982): "The Economics of Matching: Stability and Incentives," Mathematics of Operations Research, 7, 617–628.

SHORRER, R. I. (2019): "Simultaneous Search: Beyond Independent Successes," Mimeo.

SHORRER, R. I. AND S. SÓVÁGÓ (2018): "Obvious Mistakes in a Strategically Simple College Admissions Environment: Causes and Consequences," SSRN WP 2993538.

SOMERVILLE, J. (2019): "Choice-Set-Dependent Preferences in Consumer Choice: An Experimental Test," Mimeo.

UCAS (2011): "Admissions Process Review Consultation," Available at https://www.ucas.com/file/956/download?token=y8EovXLo.

FIGURE 1. Experimental Flow.

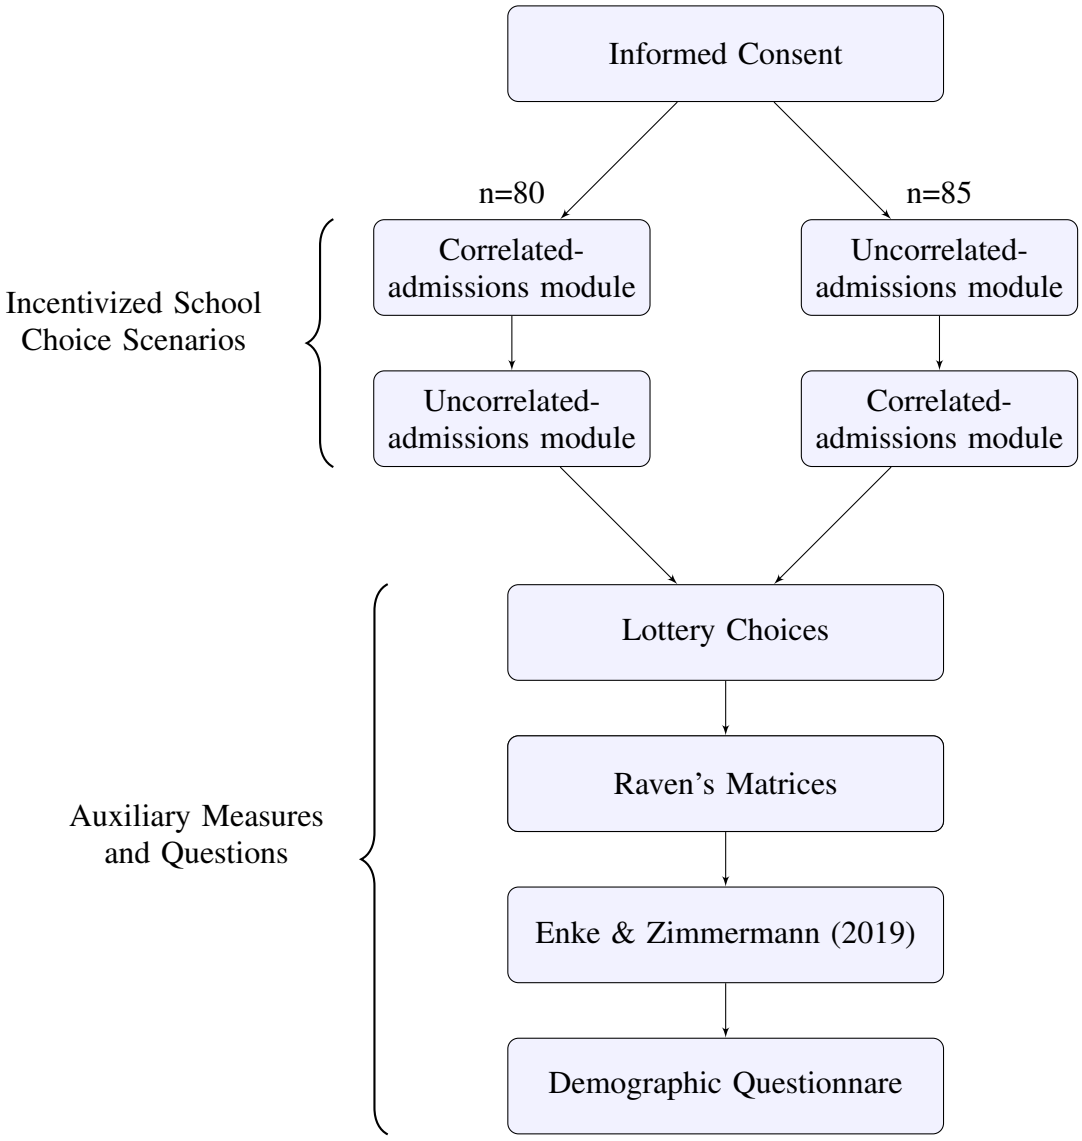


FIGURE 2. Example Set of Schools.

	Bonus if you enroll	Minimum test score
College A	$ 10	50
College B	$ 5	45
College C	$ 2.5	0

FIGURE 3. Application Strategies Across All Scenarios.

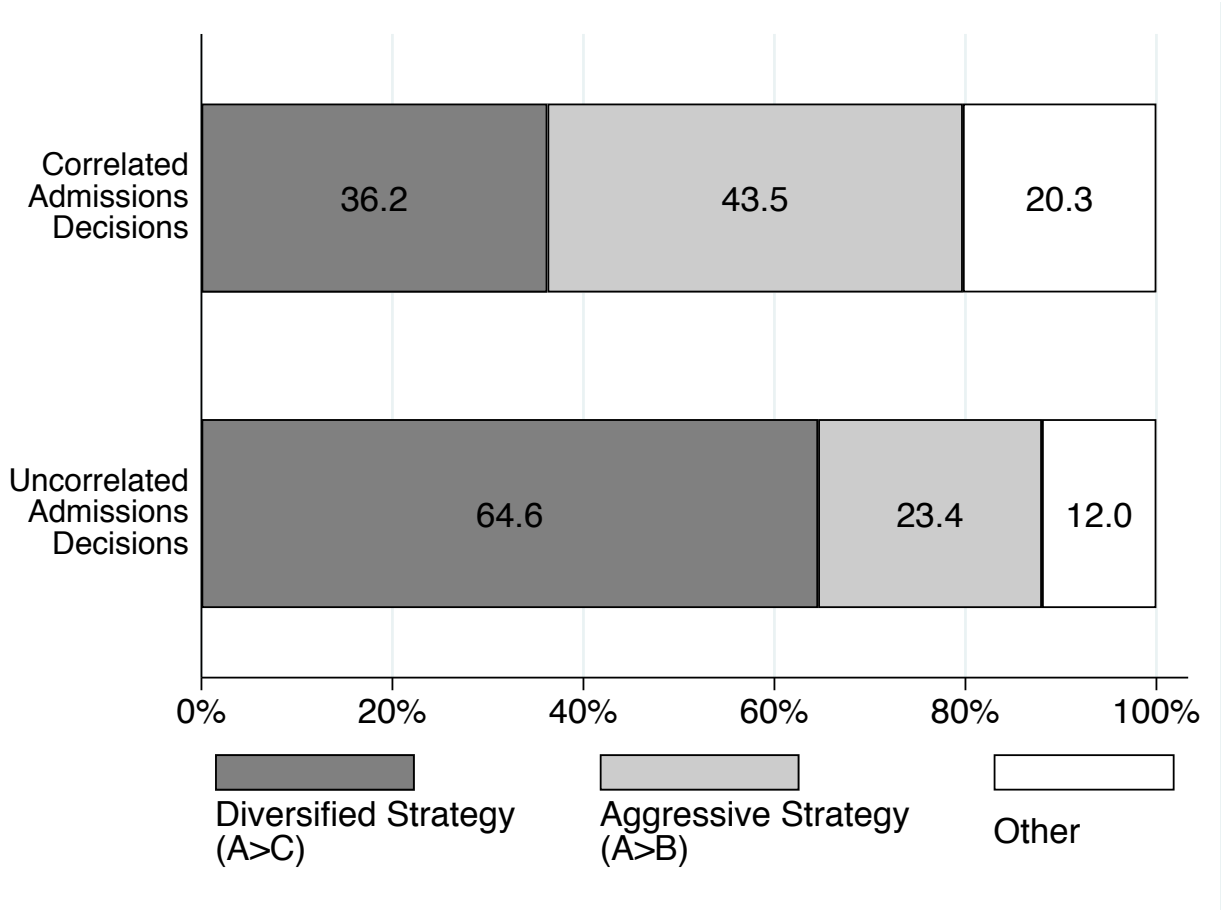


FIGURE 4. Consistency Between Application Strategies and Gamble Choices.

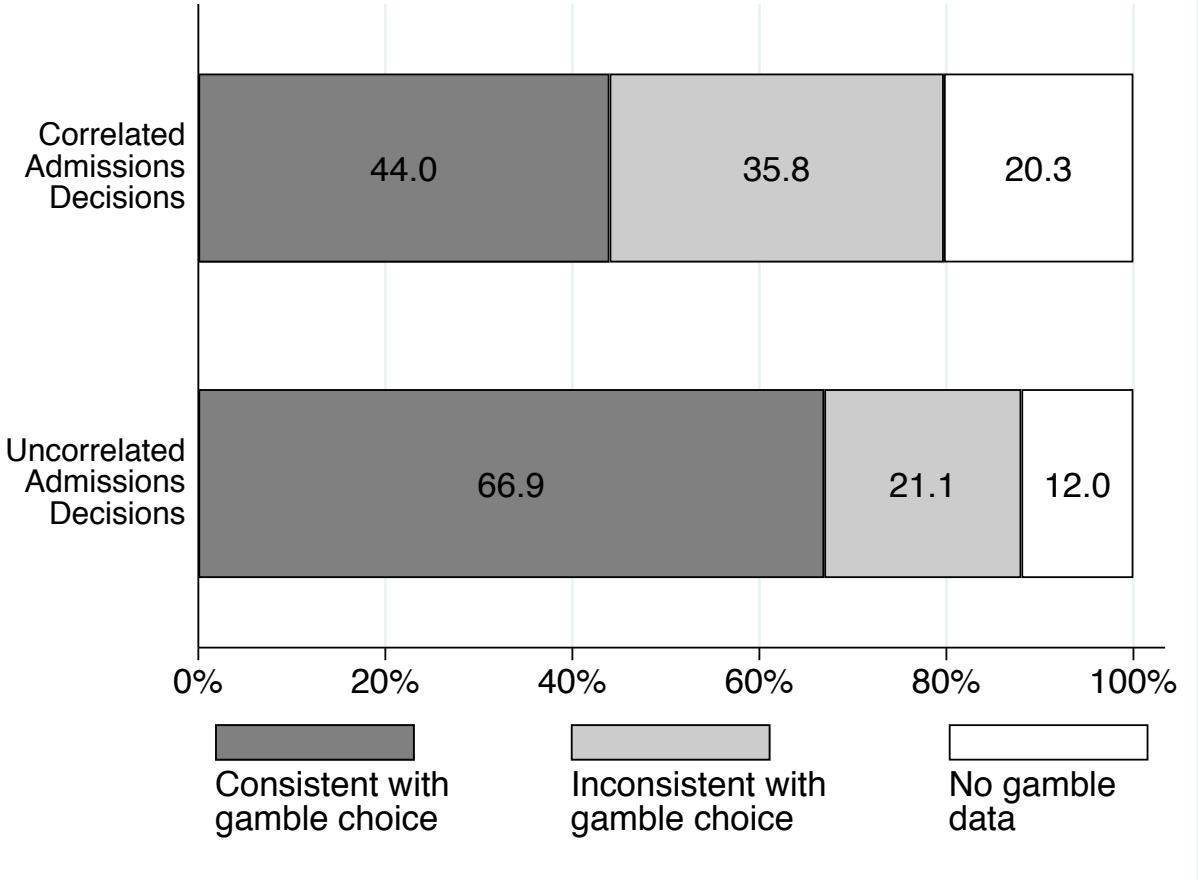
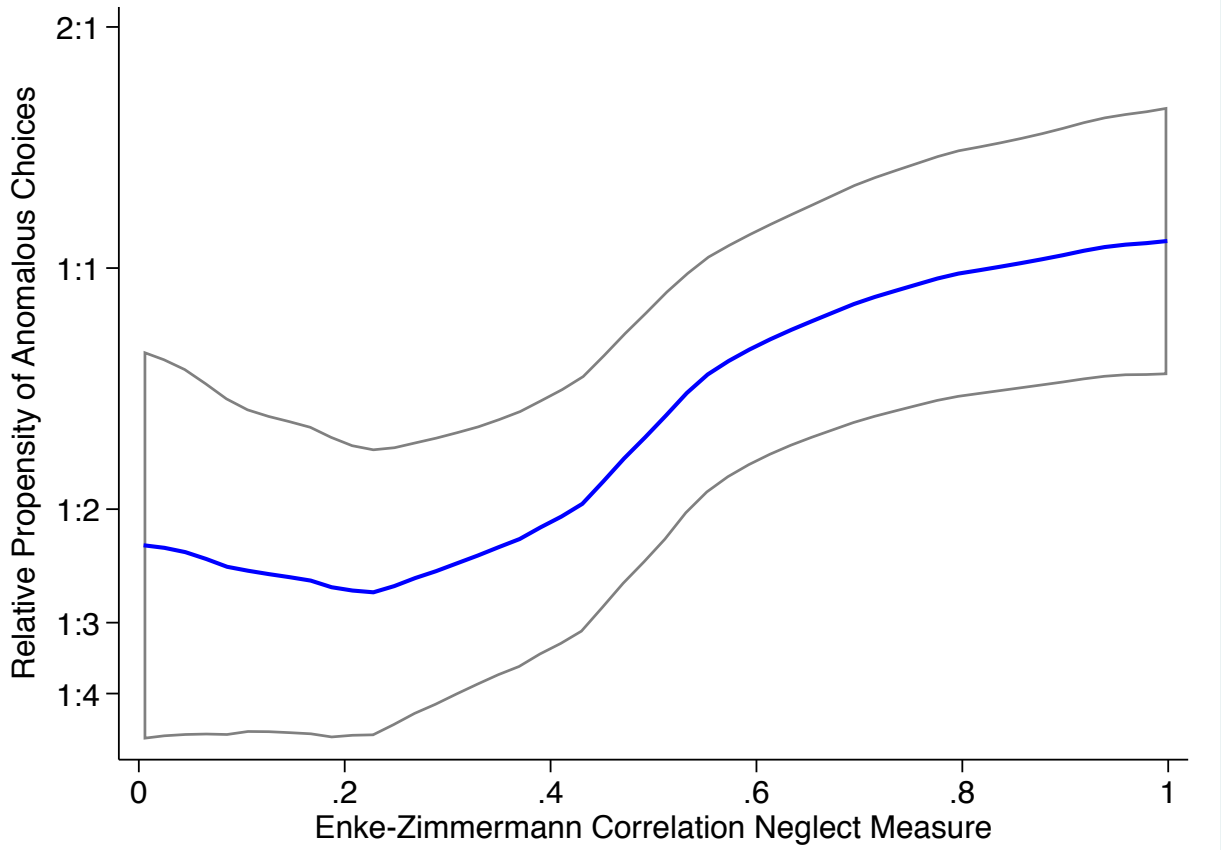


FIGURE 5. Predicting Application Mistakes with the Enke-Zimmermann Measure



Notes: This figure presents a local-polynomial regression of the relative propensity of our target preference reversal on the Enke-Zimmermann measure of correlation neglect (restricted to the unit inverval). To illustrate the interpretation of the y-axis, note that "2:1" indicates that among the nine scenarios, the subject made a correlation-neglectful preference reversal ($(A \succ B)$ under correlation and $(A \succ C)$ under independence) twice per every optimal response ($(A \succ C)$ in both framings). Bandwidth: 0.25. Kernel: Epanechnikov. Degree: 0. Confident level: 95%. Number of observations: 94.

TABLE 1. Parameters of All Scenarios.

Scenario		*Required Test Score* A	B	C	*Consequence of ROLs* (A$\succ$B)	(A$\succ$C)
1.	C	50	45	0	$\underbrace{(\$10, 0.5; \$5, 0.05; \$0, 0.45)}_{\text{EV=\$5.25}}$	$\underbrace{(\$10, 0.5; \$2.5, 0.5)}_{\text{EV=\$6.25}}$
	U	50	90	0		
2.	C	50	45	10	$\underbrace{(\$10, 0.5; \$5, 0.05; \$0, 0.45)}_{\text{EV=\$5.25}}$	$\underbrace{(\$10, 0.5; \$2.5, 0.4; \$0, 0.1)}_{\text{EV=\$6.00}}$
	U	50	90	20		
3.	C	50	20	0	$\underbrace{(\$10, 0.5; \$5, 0.3; \$0, 0.2)}_{\text{EV=\$6.5}}$	$\underbrace{(\$10, 0.5; \$2.5, 0.5)}_{\text{EV=\$6.25}}$
	U	50	40	0		
4.	C	50	20	10	$\underbrace{(\$10, 0.5; \$5, 0.3; \$0, 0.2)}_{\text{EV=\$6.5}}$	$\underbrace{(\$10, 0.5; \$2.5, 0.4; \$0, 0.1)}_{\text{EV=\$6.00}}$
	U	50	40	20		
5.	C	50	55	0	$\underbrace{(\$10, 0.5; \$0, 0.5)}_{\text{EV=\$5.00}}$	$\underbrace{(\$10, 0.5; \$2.5, 0.5)}_{\text{EV=\$6.25}}$
	U	50	100	0		
6.	C	75	60	0	$\underbrace{(\$10, 0.25; \$5, 0.15; \$0, 0.6)}_{\text{EV=\$3.25}}$	$\underbrace{(\$10, 0.25; \$2.5, 0.75)}_{\text{EV=\$4.375}}$
	U	75	80	0		
7.	C	75	60	30	$\underbrace{(\$10, 0.25; \$5, 0.15; \$0, 0.6)}_{\text{EV=\$3.25}}$	$\underbrace{(\$10, 0.25; \$2.5, 0.45; \$0, 0.3)}_{\text{EV=\$3.625}}$
	U	75	80	40		
8.	C	80	60	0	$\underbrace{(\$10, 0.2; \$5, 0.2; \$0, 0.6)}_{\text{EV=\$3.00}}$	$\underbrace{(\$10, 0.2; \$2.5, 0.8)}_{\text{EV=\$4.00}}$
	U	80	75	0		
9.	C	80	60	40	$\underbrace{(\$10, 0.2; \$5, 0.2; \$0, 0.6)}_{\text{EV=\$3.00}}$	$\underbrace{(\$10, 0.2; \$2.5, 0.4, \$0, 0.4)}_{\text{EV=\$3.00}}$
	U	80	75	50		

Notes: This table summarizes the nine "matched pairs" of scenarios presented in the experiment. Each numbered pair of rows indicates a given scenario pair. Row C presenting the test-score thresholds presented in the correlated-admissions module, whereas row U presents the test-score thresholds presented in the uncorrelated-admissions module. The last two columns of the table present the gambles induced by the two focal admissions strategies. Within the parenthesis, we present the pairs of monetary outcomes and their probabilities. Below each induced gamble, we present the expected value.

TABLE 2. Demographic Information (N=165)

	Mean	Standard Deviation	25^{th} Pctile	Median	75^{th} Pctile
Female	62.4%				
Graduated high-school in the USA	90.3%				
Degree involves Math	75.2%				
Age	21.2	3.0	20.0	21.0	22.0
High-school GPA	3.65	0.70	3.60	3.86	4.0
College GPA	3.45	0.37	3.25	3.50	3.70
Nb. Correct Raven's Matrices	5.0	1.2	4.0	5.0	6.0
Correlation-Neglect Measure	0.75	0.51	0.53	0.89	1.03

TABLE 3. Choices in the Correlated and Uncorrelated Settings.

Scenario		Rank-Order List $(A \succ B)$	$(A \succ C)$	$(B \succ C)$	Other	Test of Equality (p-values) Full Dist.	$(A \succ B)$	$(A \succ C)$
1.	C: (50, 45, 0)	48.5	44.9	4.2	2.4	0.01^{**}	$<0.01^{***}$	$<0.01^{***}$
	U: (50, 90, 0)	10.9	84.2	0.0	4.9			
2.	C: (50, 45, 10)	50.3	44.2	3.0	2.4	$<0.01^{***}$	$<0.01^{***}$	$<0.01^{***}$
	U: (50, 90, 20)	10.3	87.9	0.0	1.8			
3.	C: (50, 20, 0)	74.6	18.8	6.1	0.6	$<0.01^{***}$	$<0.01^{***}$	$<0.01^{***}$
	U: (50, 40, 0)	49.7	40.6	6.1	3.6			
4.	C: (50, 20, 10)	81.8	12.7	4.9	0.6	$<0.01^{***}$	$<0.01^{***}$	$<0.01^{***}$
	U: (50, 40, 20)	67.9	24.9	3.6	3.6			
5.	C: (50, 55, 0)	26.7	69.1	1.2	3.0	$<0.01^{***}$	$<0.01^{***}$	$<0.01^{***}$
	U: (50, 100, 0)	7.9	87.9	0.6	3.6			
6.	C: (75, 60, 0)	24.9	45.4	23.0	6.7	$<0.01^{***}$	$<0.01^{***}$	$<0.01^{***}$
	U: (75, 80, 0)	12.7	76.4	3.6	7.3			
7.	C: (75, 60, 30)	31.3	38.2	25.5	6.0	$<0.01^{***}$	$<0.01^{***}$	$<0.01^{***}$
	U: (75, 80, 40)	14.6	76.4	0.6	8.5			
8.	C: (80, 60, 0)	24.2	29.1	40	6.7	$<0.01^{***}$	0.03^{**}	$<0.01^{***}$
	U: (80, 75, 0)	14.5	57.6	18.8	9.1			
9.	C: (80, 60, 40)	30.3	23.6	39.4	6.7	$<0.01^{***}$	0.10	$<0.01^{***}$
	U: (80, 75, 50)	22.4	45.5	21.2	10.9			

Notes: This table summarizes the ROLs chosen in each matched pair of scenarios. All numbers presented (with the exception of the last column) are percentages of responses seen within a module. Columns $(A \succ B)$, $(A \succ C)$, and $(B \succ C)$ present the fractions of subjects reporting each of those ROLs, and column "other" reports the fraction of subjects reporting one of the (clearly dominated) strategies $(B \succ A)$, $(C \succ A)$, or $(C \succ B)$. The final 3 columns present p-values associated with tests for differences across the correlated and uncorrelated presentations. The column marked "Full Dist." presents the results of Fisher's exact tests of differences in the distribution of the six possible ROLs by correlation condition. The following two columns present two-sample difference-of-proportions tests, comparing the proportion picking each of the focal strategies across correlation conditions.

TABLE 4. Within-Subject Application Behavior in Matched Pairs.

Scenario	Not Influenced by Correlation		Influenced by Correlation				
	C: $(A \succ C)$ U: $(A \succ C)$	C: $(A \succ B)$ U: $(A \succ B)$	C: $(A \succ B)$ U: $(A \succ C)$	C: $(A \succ C)$ U: $(A \succ B)$	Included $(B \succ C)$	Inc. Reverse Alphabetical	Frac choosing U: $(A \succ C)$ given C: $(A \succ B)$.
1. C: (50, 45, 0) U: (50, 90, 0)	41.2	8.5	37.6	2.4	4.2	5.9	77.5
2. C: (50, 45, 10) U: (50, 90, 20)	43.6	7.9	41.2	0.6	1.8	4.9	81.9
3. C: (50, 20, 0) U: (50, 40, 0)	12.7	42.4	25.5	6.0	9.7	3.7	34.1
4. C: (50, 20, 10) U: (50, 40, 20)	7.9	60.6	14.6	4.9	7.9	4.1	17.8
5. C: (50, 55, 0) U: (50, 100, 0)	64.9	4.9	21.2	3.0	1.2	4.8	79.5
6. C: (75, 60, 0) U: (75, 80, 0)	41.8	9.7	13.3	1.8	21.8	11.6	53.7
7. C: (75, 60, 30) U: (75, 80, 40)	33.9	10.9	17.6	3.0	22.4	12.7	58.0
8. C: (80, 60, 0) U: (80, 75, 0)	22.4	10.3	8.5	1.8	45.5	11.5	35.0
9. C: (80, 60, 40) U: (80, 75, 50)	17.6	13.9	10.9	2.4	40.6	14.6	36.0
C: Average U: Average	31.8	18.8	21.2	2.9	17.2	8.2	52.6

Notes: This table summarizes the within-subject patterns of ROLs submitted for each matched pair of scenarios. We focus on characterizing the use of the two focal application strategies, the aggressive application strategy ($(A \succ B)$) and the diversified application strategy ($(A \succ C)$). The first two columns present the fraction of subjects pursuing the diversified or the aggressive strategies, respectively, under both correlation conditions. These are the subjects who pursue a focal strategy and are not influenced by correlation. The next two columns present the fraction of subjects who switch from one focal strategy to the other based on the correlation condition—i.e., the subjects who are influenced by correlation. The next two columns characterize the fraction of subjects for whom at last one submitted ROL was a non-focal strategy: either $(B \succ C)$ or one in which programs are submitted in reverse-alphabetical order. The final column presents the fraction of subjects who chose the diversified strategy in the uncorrelated-admissions module conditional on choosing the aggressive strategy in the correlated-admissions module.

TABLE 5. Choices in Lotteries.

	Lottery question	% chose ($A \succ C$) in lottery	% chose ($A \succ B$) in lottery cond. on ROL responding to correlation	% chose ($A \succ C$) in lottery cond. on ($A \succ C$) in both ROLs
1.	($A \succ B$) : \$10 w/50%, \$5 w/5% ($A \succ C$): \$10 w/50%, \$2.5 w/50%	92.7	9.7	92.7
2.	($A \succ B$) : \$10 w/50%, \$5 w/5% ($A \succ C$): \$10 w/50%, \$2.5 w/40%	97.6	1.5	98.6
3.	($A \succ B$) : \$10 w/50%, \$5 w/30% ($A \succ C$): \$10 w/50%, \$2.5 w/50%	68.5	21.4	76.2
4.	($A \succ B$) : \$10 w/50%, \$5 w/30% ($A \succ C$): \$10 w/50%, \$2.5 w/40%	47.3	62.5	69.2
5.	($A \succ B$) : \$10 w/50% ($A \succ C$): \$10 w/50%, \$2.5 w/50%	97.0	5.7	97.2
6.	($A \succ B$) : \$10 w/25%, \$5 w/15% ($A \succ C$): \$10 w/25%, \$2.5 w/75%	98.2	0.0	98.6
7.	($A \succ B$) : \$10 w/25%, \$5 w/15% ($A \succ C$): \$10 w/25%, \$2.5 w/45%	98.2	6.9	100
8.	($A \succ B$) : \$10 w/20%, \$5 w/20% ($A \succ C$): \$10 w/20%, \$2.5 w/80%	97.0	14.3	97.3
9.	($A \succ B$) : \$10 w/20%, \$5 w/20% ($A \succ C$): \$10 w/20%, \$2.5 w/40%	88.5	38.9	100
	C: Average U: Average	87.2	17.9	92.2

Notes: This table summarizes the choices made over pairs of gambles constructed to offer the same payouts as arise from each scenario's focal strategies. The first column reports the fraction of subjects choosing the gamble that arises from the diversified application strategy, illustrating that this option is overwhelmingly preferred when the consequences are made transparent. The second column shows the fraction of subjects choosing the gamble that arises from the aggressive application strategy contingent on being coded as responding to correlation in the analysis of Table 4. The third column shows the fraction of subjects choosing the gamble associated with the diversified application strategy conditional on pursuing that strategy in both correlation conditions.

TABLE 6. Predicting Correlation-Neglectful Preference Reversals.

	(1)	(2)
Enke-Zimmermann Measure	0.323 (0.120)***	0.263 (0.125)**
EZ Missing	0.212 (0.100)**	0.171 (0.105)
Raven's Matrices Performance		-0.046 (0.025)*
Female		0.036 (0.060)
High School GPA		0.007 (0.042)
College GPA		-0.090 (0.082)
Attended High School in USA		-0.043 (0.099)
Math		0.052 (0.073)
Constant	0.220 (0.090)**	0.751 (0.331)
# of observations	157	157
R^2	0.045	0.080

Notes: This table presents OLS regressions of our measure of the rate of correlation-neglectful preference reversals on the Enke-Zimmermann measure of correlation neglect. The Enke-Zimmermann measure is treated as missing if it is measured outside of the unit interval, in which case the variable "EZ Missing" is set to 1. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Appendices

A. Proofs

Proposition 1. *For any school choice environment and for any undominated ROL r,* $V_s(r) \leq V_n(r)$.

Proof. The proof proceeds by induction on the size of the ROL. The case $k = 1$ is obvious as correlation in admission decisions is irrelevant for students' subjective expected utility. For the case $k > 1$, let $r^{2:k}$ denote the continuation ROL from the second to the k-th ranked schools. Then, we have that for all $x \in \{s, n\}$,

$$V_x(r) = (1 - F(c_{r^1}))\, u\left(r^1\right) + F(c_{r^1})\, V_x(r^{2:k} \mid \text{rejected by } r^1).$$

For the neglectful type, $V_n(r^{2:k} \mid \text{rejected by } r^1) = V_n(r^{2:k})$. For the sophisticated type, $V_s(r^{2:k} \mid \text{rejected by } r^1) \leq V_s(r^{2:k})$, as the absence of information results in a first order stochastically higher distribution of outcomes (mass is reduced proportionally from all options and added to r^2).[24]

By induction, $V_n(r^{2:k}) \geq V_s(r^{2:k})$. Altogether we have that

$$V_n(r^{2:k} \mid \text{rejected by } r^1) = V_n(r^{2:k}) \geq V_s(r^{2:k}) \geq V_s(r^{2:k} \mid \text{rejected by } r^1),$$

and hence

$$V_n(r) \geq V_s(r).$$

□

Proposition 2. *For any integer k, and any decision environment where the agent is constrained to (costlessly) apply to up-to-k schools, the price of neglect for the neglectful type is bounded above*

[24] Here we use the fact that undominated ROLs are ordered according to true preferences.

by $1-\frac{1}{k}$. *Furthermore, this bound is tight—for any* k*, there exist school choice environments where the price of neglect is arbitrarily close to* $1-\frac{1}{k}$.

Proof. To begin with, note that the optimal size-1 ROL is identical for all types, as correlation only matters when applying to multiple schools. Next, observe that the neglectful type believes that admissions decisions across schools are independent. Thus, by Theorem 1 of Chade and Smith (2006), any subjective-optimal ROL of size-k of the neglectful type includes a subjective-optimal singleton ROL, which is also an objective optimal singleton ROL by the first observation. Thus, subjective-optimal ROLs of size-k achieve at least as much experienced utility as the optimal size-1 ROL (the fact that the ROL of the neglectful type includes more schools can only improve the utility he will experience, as he will attend the best school that accepts him). Finally, by Theorem 2 of Shorrer (2019), the expected utility of a sophisticated agent from an optimal size-1 ROL is greater than or equal to $\frac{1}{k}$ of the expected utility from an optimal size-k ROL.

To see that the lower bound is tight, consider an arbitrarily small $\epsilon > 0$. For $m \in \{1,2,...,k-1\}$, let $u_m := \epsilon^{-m}$ and let $c_m := 1-\epsilon^m$, and let $u_k := \epsilon^{-k}(1+\delta)$ and $c_k := 1-\epsilon^k$. Let X consists of k copies of each type of school, (u_i, c_i). Then the full correlation neglectful type will choose the k copies of the most desired school, u_k, and get utility of $1+\delta$ (see, e.g. Chade and Smith 2006). But by choosing one school of each type the expected utility is approximately k for sufficiently small ϵ. □

Proposition 3. *For any constraint on the size of the ROL* k*, the neglectful type is at least as likely to be unassigned as the sophisticated type.*

Our leading example shows that this comparison may be strict.

Proof. The case of $k = 1$ is obvious since correlation plays no role when students can only apply to one school. Next, recall from Shorrer (2019) that options that are more selective and less desirable than other options are dominated and do not appear on an optimal ROL of a sophisticated agent.

A consequence of this statement is that when considering the sophisticated type's ROL, there is no loss in focusing on a subset of undominated alternatives $X' \subset X$ such that for any $x, y \in X'$, $u(x) > u(y) \iff c_x > c_y$.

Consider the subjective-optimal size-k ROL of the neglectful type, $r_n(k)$. Since options in $X \backslash X'$ are dominated, they can only appear on ROLs that include choices that dominate them. Thus, the least selective school on $r_n(k)$ belongs to X'. Hence,$|r_n(k) \cap (X \backslash X')| = m < k$.

Consider the choice problem where an agent need to choose optimal ROLs of size $k - m$ from X' with the stochastic outside option of m independent lotteries, one for each $i \in r_n(k) \cap (X \backslash X')$, where the probability of realization of lottery i is $1 - c_i$ and its utility from attending is $u(i)$ (but the student still can only attend one school). The outside option is how the neglectful type subjectively perceives $r_n(k) \cap (X \backslash X')$. Note that since optimal ROLs rank schools according to desirability, in this problem, the lowest-ranked school is the least selective option in the ROL (ignoring the outside options).

We now claim that the last (i.e., $k - m$-th ranked) school on the neglectful type's ROL is associated with a weakly more selective cutoff than that of the sophisticated type. Towards contradiction, assume the opposite. Then the last choice on the neglectful type's ROL is less selective and thus less desirable than the sophisticated type's last choice (since choices are from X'). This means that the choice does not appear on the ROL of the sophisticated type, thus he can deviate and replace his last choice with the neglectful type's last choice. But because both agents choose their last school conditional on rejection by all higher ranked schools (Shorrer 2019, Lemma 2), the sophisticated type's beliefs are MLRP-lower and this is a contradiction to Proposition 2 of Shorrer (2019), which states that sophisticated agents with higher beliefs apply more aggressively (as the lack of sophistication does not play a role in ROLs of size 1 – except, of course, for the effect of false beliefs from conditioning on previous rejections).

Next, note that the neglectful type's subjective optimal ROL must be identical to his optimal ROL in the original problem (the outside option and the constraint on the size of the ROL were chosen to mimic a situation where his strategy space was restricted to include certain options which appear on his subjective optimal ROL anyway and to not include certain school that did not appear on his subjective optimal ROL anyway).

Lastly, note that if we remove the sophisticated type's access to the outside option, his ROL becomes less aggressive (Shorrer 2019, Theorem 3). And the least selective school on his ROL of size k in this problem is even less selective than the least selective school on his optimal ROL of size $k-m$ (Shorrer 2019, Theorem 1). But the optimal ROL of size k from X' coincides with the optimal ROL of size k from X (Shorrer 2019, Lemma 1). Together with the fact that the least selective school on $r_n(k)$ belongs to X', this completes the proof. □

B. SUPPLEMENTAL ANALYSES

TABLE A1. Choices in the correlated and uncorrelated settings by order of module.

Scenario		Uncorrelated Module First $(A \succ B)$	$(A \succ C)$	$(B \succ C)$	Other	Correlated Module First $(A \succ B)$	$(A \succ C)$	$(B \succ C)$	Other	Test of Equality p-value
1.	C: (50, 45, 0)	45.9	48.2	4.7	1.2	51.3	41.3	3.8	3.7	0.66
	U: (50, 90, 0)	8.2	87.1	0.0	4.7	13.8	81.3	0	5.0	0.65
2.	C: (50, 45, 10)	52.9	42.4	3.5	1.2	47.5	46.3	2.5	3.8	0.80
	U: (50, 90, 20)	9.4	90.6	0.0	0.0	11.3	85.0	0.0	3.8	0.22
3.	C: (50, 20, 0)	72.9	21.1	5.9	0.0	76.3	16.3	6.2	1.2	0.71
	U: (50, 40, 0)	52.9	38.8	5.9	2.4	46.3	42.5	6.3	5.0	0.89
4.	C: (50, 20, 10)	78.8	12.9	7.1	1.2	85.0	12.5	2.5	0.0	0.44
	U: (50, 40, 20)	75.3	20.0	1.2	3.5	60.0	30.0	6.3	3.8	0.02**
5.	C: (50, 55, 0)	21.2	76.5	1.2	1.2	32.5	61.3	1.3	5.0	0.14
	U: (50, 100, 0)	9.4	87.1	1.2	2.4	6.3	88.8	0.0	5.0	0.52
6.	C: (75, 60, 0)	20.0	50.6	25.9	3.5	30.0	40.0	20.0	10.0	0.20
	U: (75, 80, 0)	12.9	77.7	1.2	8.2	12.5	75.0	6.3	6.3	0.24
7.	C: (75, 60, 30)	28.2	42.4	22.4	3.7	32.5	33.8	28.8	5.0	0.76
	U: (75, 80, 40)	14.1	77.7	1.2	7.2	15.0	75.0	0.0	10.0	0.97
8.	C: (80, 60, 0)	21.2	34.1	38.8	5.9	27.5	23.8	41.3	7.5	0.39
	U: (80, 75, 0)	15.3	55.3	20.0	9.4	13.8	60.0	17.5	8.8	0.95
9.	C: (80, 60, 40)	31.8	25.9	40.0	2.4	28.8	21.3	38.8	11.2	0.19
	U: (80, 75, 50)	25.9	41.2	21.2	11.8	18.8	50.0	21.3	9.5	0.73

Notes: This table summarizes the ROLs chosen in each matched pair of scenarios by which module subjets saw first. All numbers presented (with the exception of the last column) are percentages of responses seen within a module. Columns $(A \succ B)$, $(A \succ C)$, and $(B \succ C)$ present the fractions of subjects reporting each of those ROLs, and column "other" reports the fraction of subjects reporting one of the (clearly dominated) strategies $(B \succ A)$, $(C \succ A)$, or $(C \succ B)$. The final column present the p-value associated with Fisher's exact test for differences across populations who saw the uncorrelated or correlated module, using the full distribution of choices without aggregating dominated ROLs.

TABLE A2. Choices in the Correlated and Uncorrelated Settings (using first module only).

		Rank-Order List				Test of Equality (p-values)		
	Scenario	$(A \succ B)$	$(A \succ C)$	$(B \succ C)$	Other	Full Dist.	$(A \succ B)$	$(A \succ C)$
1.	C: (50, 45, 0)	51.3	41.3	3.8	3.7	<0.01***	<0.01***	<0.01***
	U: (50, 90, 0)	8.2	87.1	0.0	4.7			
2.	C: (50, 45, 10)	47.5	46.3	2.5	3.8	<0.01***	<0.01***	<0.01***
	U: (50, 90, 20)	9.4	90.6	0.0	0.0			
3.	C: (50, 20, 0)	76.3	16.3	6.2	1.2	<0.01***	<0.01***	<0.01***
	U: (50, 40, 0)	52.9	38.8	5.9	2.4			
4.	C: (50, 20, 10)	85.0	12.5	2.5	0.0	0.16	0.12	0.19
	U: (50, 40, 20)	75.3	20.0	1.2	3.5			
5.	C: (50, 55, 0)	32.5	61.3	1.3	5.0	<0.01***	<0.01***	<0.01***
	U: (50, 100, 0)	9.4	87.1	1.2	2.4			
6.	C: (75, 60, 0)	30.0	40.0	20.0	10.0	<0.01***	<0.01***	<0.01***
	U: (75, 80, 0)	12.9	77.7	1.2	8.2			
7.	C: (75, 60, 30)	32.5	33.8	28.8	5.0	<0.01***	<0.01***	<0.01***
	U: (75, 80, 40)	14.1	77.7	1.2	7.2			
8.	C: (80, 60, 0)	27.5	23.8	41.3	7.5	<0.01***	0.06*	<0.01***
	U: (80, 75, 0)	15.3	55.3	20.0	9.4			
9.	C: (80, 60, 40)	28.8	21.3	38.8	11.2	0.03**	0.68	<0.01***
	U: (80, 75, 50)	25.9	41.2	21.2	11.8			

Notes: This table summarizes the ROLs chosen in each matched pair of scenarios. All numbers presented (with the exception of the last column) are percentages of responses seen within a module using the first module of each treatment only. Columns $(A \succ B)$, $(A \succ C)$, and $(B \succ C)$ present the fractions of subjects reporting each of those ROLs, and column "other" reports the fraction of subjects reporting one of the (clearly dominated) strategies $(B \succ A)$, $(C \succ A)$, or $(C \succ B)$. The final 3 columns present p-values associated with tests for differences across the correlated and uncorrelated presentations. The column marked "Full Dist." presents the results of Fisher's exact tests of differences in the distribution of the six possible ROLs by correlation condition. The following two columns present two-sample difference-of-proportions tests, comparing the proportion picking each of the focal strategies across correlation conditions.