NBER WORKING PAPER SERIES

THE MEMORY PREMIUM

Yuval Salant
ⓡ
Jorg L. Spenkuch
ⓡ
David Almog

Working Paper 33649
http://www.nber.org/papers/w33649

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2025, Revised August 2025

The first two authors are listed in random order. We gratefully acknowledge helpful comments from Ido Ben Artzi, Ariel Rubinstein, David Smerdon, and Leeat Yariv, as well as from seminar audiences at Penn, Penn State, Toronto, USC, and the Nash 75 Conference at Oxford. We are also grateful to Thibault Duplessis and Theo Wait for providing us with data from Lichess, and for patiently answering our questions. Michael Indenkämpen provided excellent research assistance. The research in this paper was supported in part through the computational resources provided by the Quest high-performance computing facility at Northwestern University. The experiment in this paper has been reviewed by the Northwestern IRB, and was preregistered in the OSF Registries at https://osf.io/8tkd4. Correspondence should be addressed to the first two authors at: MEDS Department, Kellogg School of Management, 2211 Campus Dr, Evanston, IL 60208; or by email to y-salant@kellogg.northwestern.edu and j-spenkuch@kellogg.northwestern.edu. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2025 by Yuval Salant ⓡ Jorg L. Spenkuch ⓡ David Almog. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Memory Premium
Yuval Salant ⓡ Jorg L. Spenkuch ⓡ David Almog
NBER Working Paper No. 33649
April 2025, Revised August 2025
JEL No. D01, D03, D87, D91

ABSTRACT

We explore the role of memory for choice behavior in unfamiliar environments. Drawing on a unique data set, we document that decision-makers exhibit a "memory premium." They tend to choose in-memory alternatives over out-of-memory ones, even when the latter are objectively better. Consistent with well-established regularities regarding the inner workings of human memory, the memory premium is associative, subject to interference and repetition effects, and decays over time. Evidence from a laboratory experiment implies that the memory premium arises because individuals rely on their memory to identify alternatives that warrant careful consideration, albeit in a boundedly rational way.

Yuval Salant
Northwestern University
Kellogg School of Management
y-salant@kellogg.northwestern.edu

ⓡ

Jorg L. Spenkuch
Northwestern University
Kellogg School of Management
and NBER
j-spenkuch@kellogg.northwestern.edu

ⓡ

David Almog
Northwestern University
david.almog@kellogg.northwestern.edu

1. Introduction

The goal of this paper is to better understand how memory retrieval affects choice behavior. Cognitive scientists generally conceptualize memory as a three-stage process: encoding, storage, and retrieval (Melton 1963; Baddeley et al. 2009). In the encoding stage, the brain transforms sensory inputs into neural representations. These representations are consolidated and stored over time before being accessed and utilized during the retrieval stage.

Our first contribution is to document that, in unfamiliar environments, decision-makers are more likely to choose alternatives that are retrievable from memory than those that are not, even when the latter are objectively better. This phenomenon, which we call the "memory premium," can be understood as an extension of the availability heuristic to the domain of choice. In the words of Tversky and Kahneman (1973, p. 208), "a person is said to employ the availability heuristic whenever he estimates frequency and probability by the ease with which instances or associations could be brought to mind." Whereas the availability heuristic concerns judgments and probability assessments, the memory premium connects recall and choice behavior.

Our second contribution is to develop a set of testable predictions that relate the basic cognitive mechanisms governing memory retrieval to choice behavior. Previous research in psychology and neuroscience establishes four stylized facts about the retrieval process (see, e.g., Kahana 2012): (i) Memory is associative. Mental representations are interconnected, and past experiences are more easily recalled when they are related to current environmental cues. (ii) Memory decays over time. Recent experiences are more likely to be retrievable than temporally distant ones. (iii) Memory is strengthened by repetition. Experiences that are encountered more frequently are more likely to be recalled than single-instance encounters. (iv) Retrieval is subject to interference. Newer memories disrupt the recall of older ones, and vice versa. Taking these fundamental mechanisms as given, we document that the memory premium exhibits analogous comparative statics. It is associative, subject to repetition and interference effects, and it decays over time. As a result, the cognitive mechanisms behind recall also predict choices.

Our findings are based on a unique observational data set on choice behavior in unfamiliar environments in conjunction with a tightly controlled laboratory experiment. In our observational data, we observe the opening moves of nearly one hundred and fifty thousand experienced chess players as they engage in a game of Chess960 for the first time. Chess960 is a variant of classical chess. It uses the same board and pieces, and it relies on the same rules. The key difference relative to regular chess is that the initial positions of the pieces on players' back ranks are randomized, resulting in 960 different possible starting positions.

Chess960 provides an almost ideal setting to study the influence of memory on choice

behavior in unfamiliar environments. In regular chess, players typically invest significant amounts of time and effort into memorizing effective opening moves and counter-strategies. In Chess960, however, the large number of potential starting positions makes rote memorization of opening sequences impractical. Moreover, given that even minor differences in the initial placement of pieces can have profound implications for optimal strategy, opening moves that tend to perform well in classical chess need not be a good choice in a particular starting position in Chess960. As a consequence, players should rely on their analytical skills rather than memory when selecting an opening move. In our analysis, we ask to which extent they nonetheless draw on memory and how this reliance depends on features of the new environment.

Our data come from `lichess.org`, one of the most popular internet chess servers. Every day, Lichess hosts several million games of classical chess and thousands of games of Chess960.[1] We have information on all moves in the universe of games on the platform from January 2013 through June 2021.

There are two key challenges to empirically documenting the memory premium in a real-world environment. The first one is to carefully control for the quality of the available alternatives. Accounting for quality is important in order to rule out that in-memory alternatives are chosen more frequently because they are better than out-of-memory ones. In the context of Chess960, we can overcome this difficulty by using sophisticated computer algorithms to objectively measure the quality of all available opening moves. These algorithms have proven to be considerably better at analyzing chess positions than even the best grandmasters.

The second challenge is to approximate players' (unobservable) memory. Here, we draw on the richness of our data. Since we observe each player's history of opening moves in all standard chess games on the platform, we can model her memory as a function of this history and the corresponding outcomes prior to her first game of Chess960. According to our workhorse definition, a given opening move is "in-memory" for a particular player if she previously used that move in at least one standard chess game. This is a conservative definition because, due to decay, not every memory of a move that should, in principle, be retrievable can, in fact, be recalled. We also explore alternative definitions of memory, such as classifying a move as in-memory if it had been played sufficiently often or within a certain period of time.

Irrespective of particulars, we find that, in their first game of Chess960, players are significantly more likely to choose an in-memory opening move than an equally good out-of-memory alternative. Relying on our workhorse definition of memory, the average memory

[1] Lichess also allows users to play other variants of chess, most of which are more obscure than Chess960.

premium equals 3.8–4.5 percentage points (p.p.), or about 75–89% of the mean choice frequency in this setting. For comparison, taking our regression estimates at face value, the implied difference in the choice frequency of moves at the fifth and ninety-fifth percentile of the quality distribution is only about 3.2 p.p. The average memory premium is, therefore, large enough to overcome even significant differences in the objective quality of moves.

The average memory premium, however, masks considerable heterogeneity. The memory premium of inexperienced players is about twice as large as that of experienced ones, and "good" opening moves exhibit larger premia than "bad" ones. Moreover, opening moves that are common in standard chess exhibit much higher premia than moves that are only rarely used. Importantly though, we detect an economically meaningful memory premium for all opening moves, even in starting positions in which they are not close to being optimal.

In order to link the memory premium to the inner workings of human memory, we develop comparative statics that bridge the gap between well-established cognitive mechanisms and choice behavior. For example, based on associative recall, we predict that the memory premium increases in the similarity between the starting position in Chess960 and the standard chess board. Measuring similarity as the number of chess pieces that are placed on the same square, this is indeed what we find. Overall, we provide strong evidence that choice behavior is predicted by the same cognitive forces that have previously been shown to influence recall.

Our main results focus on players' first game of Chess960. We, therefore, document the existence of the memory premium in an unfamiliar choice environment. Does the memory premium decrease as individuals become more familiar with the setting? In order to speak to this question, we expand the scope of our analysis to all opening moves in players' first fifty games of Chess960 as White. As expected, the memory premium shrinks considerably as individuals gain experience; but it does not vanish. In fact, after about twenty-five games, the estimated memory premium stabilizes around 2.5 p.p., or about 50% of the average choice frequency.

There are at least two classes of mechanisms that may give rise to the memory premium. The first one involves maximization coupled with subjective differences in the expected utility of in- and out-of-memory alternatives. According to this explanation, individuals assign higher expected utility to the former, even when the latter are objectively just as good. Such differences in expected utility may arise because individuals extrapolate from the familiar environment in which in-memory alternatives perform well, or because decision-makers are risk-averse and hold more precise beliefs about alternatives that they have previously chosen. The second class of mechanisms involves deviations from utility maximization. Rather than considering all of the available options, decision-makers may rely on choice procedures that prioritize in- over out-of-memory alternatives. Examples of such procedures include

maximization from consideration sets that disproportionately include in-memory options, satisficing with a search order that prioritizes in-memory alternatives, and other choice processes whereby alternatives that are retrievable from memory are more likely to be seriously considered.

To better understand the mechanism behind the memory premium and to replicate this hitherto unknown phenomenon in a different setting, we conducted a pre-registered laboratory experiment. In the experiment, subjects are asked to choose between six investment options, half of which are in-memory on account of having been encountered earlier in the experiment. By randomly assigning memory status and by *inducing* values and beliefs, our experimental design eliminates any systematic differences in the subjective utility of in- and out-of-memory alternatives. Crucially, the design makes observable which alternatives participants consider and in what order.

Our experimental results point to a procedural explanation for the memory premium. In-memory options are not only more likely to be chosen than out-of-memory ones, but they are more likely to be considered for choice, and are more likely to be considered first—even when the most promising in-memory alternative is worse than the average out-of-memory option. More broadly, our findings suggest that memory shapes choice behavior because individuals draw on their memory in order to identify alternatives that warrant closer examination.

We conclude by proposing that relying on memory to guide decision-making in unfamiliar environments may be boundedly rational; it economizes on cognitive costs while yielding favorable outcomes—provided the new environment is sufficiently similar to the familiar one. We illustrate this point by drawing directly on our observational data. Since we observe the objective value of all available opening moves in Chess960, we can ask how the quality of players' choices would vary under different choice procedures. Our counterfactuals establish that a decision-maker could, in expectation, improve upon random choice by randomly selecting an in-memory alternative; and the magnitude of the improvement would increase in the degree of similarity between the starting position in Chess960 and that in standard chess. Our counterfactuals also consider decision makers who maximize from consideration sets (Manzini and Mariotti 2014). We show that, except in highly dissimilar board positions, restricting the consideration set to in-memory moves would cost such a decision maker very little. Analyzing players' *actual* choices, we observe that, on average, they do better when they choose in- rather than out-of-memory moves. Only in the most dissimilar starting positions do in-memory choices tend to be worse.

Related Literature Our findings contribute to a growing literature on the role of memory in economic decision-making. An early decision-theoretic contribution to this literature is

Gilboa and Schmeidler (1995, 2001). In their theory of case-based decisions, the value of an alternative is a weighted sum of its past performance, with larger weights attached to outcomes in decision problems that are more similar to the current one. In another early contribution, Mullainathan (2002) develops a model in which associative recall and rehearsal affect beliefs about a state variable, such as income. More recently, Bordalo et al. (2020, 2022, 2024) incorporate associative memory and interference into models of belief formation. For example, Bordalo et al. (2020) present a theory of consumer choice with selective recall. In their model, the decision maker relies on memories of past purchases to form a norm, or expectation, about the quality and price of the good in the current environment, which in turn influences purchase decisions. Other important contributions include Azeredo da Silveira et al. (2019; 2024), Nagel and Xu (2022), and Wachter and Kahana (2024).

Unlike extant theoretical work, our investigation does not micro-found how decision makers use memory to evaluate an object or form expectations about a payoff-relevant parameter. Instead, we examine the effect of memory on the mapping from choice situations to outcomes. As we explain in Section 2, there are several natural choice procedures that can give rise to the memory premium, including maximization with memory-dependent perceptions of utility, maximization from consideration sets, and satisficing.

There is also an empirical literature on the effect of memory on belief formation and decision-making (see, e.g., Afrouzi et al. 2023; Charles 2022, 2024; D'Acunto et al. 2021; Enke et al. 2024; Hartzmark et al. 2021; Jiang et al. 2024). In influential work, Malmendier and Nagel (2011, 2016) show that investors who lived through the Great Depression are less likely to invest in the stock market later in life, and that experiencing high inflation has lasting effects on inflation expectations.[2] Graeber et al. (2024) show that selective memory explains why the effect of stories on beliefs decays less strongly than the effect of statistics. We contribute to this empirical literature by identifying a novel feature of memory-based choice behavior. Our findings thus help to shed light on the cognitive mechanisms underlying real-world decision-making.

2. Conceptual Framework

To guide our empirical analysis, we consider the mapping from choice situations to choice probabilities, conditional on the memory of the decision maker (DM). We propose a new feature of this mapping, which we call the memory premium; and we formulate testable comparative statics that the memory premium should satisfy.

[2]See also Malmendier and Wachter (2024) for a survey of the literature on how past experiences shape expectations and behavior.

2.1. *Preliminaries*

Fix a grand set of alternatives $\mathcal{X}$ and a collection of environments $\mathcal{E}$. Alternatives have objective value to the DM, which may depend on the environment, and which the DM may or may not know. We use $V(x, E)$ to denote the objective value of alternative $x \in \mathcal{X}$ in environment $E \in \mathcal{E}$. Environments may be more or less similar to one another, with $\sigma(E_1, E_2)$ denoting the degree of similarity between environments E_1 and E_2.

A choice situation is a pair (A, E), where $A \subset \mathcal{X}$ includes all the alternatives available for choice, and E is the environment in which choices are to be made. The DM encountered a fixed choice situation $\hat{S} = (\hat{A}, \hat{E})$ in the past. We refer to $\hat{E}$ as the "familiar" environment and to all other environments as unfamiliar.

The DM's memory M consists of all experiences she had in the familiar environment. For simplicity, we equate an experience with a chosen alternative. That is, $M = (a(1), a(2), ...)$, where $a(k) \in \hat{A}$ is the alternative that the DM chose k periods ago from $\hat{S}$.[3] We refer to an alternative a as "in memory" and write $a \in M$ if there exists a k such that $a = a(k)$. An alternative a that is not in memory M is said to be "out of memory." In this case, we write $a \notin M$. Given a memory M and an in-memory alternative a, we denote by M^{-a} the memory that is obtained by removing all instances of a from M.

2.2. *Memory Premium*

We are interested in studying memory-related properties of the random choice function $C(S, a \mid M)$. This function assigns, to every choice situation $S = (A, E)$ and any alternative $a \in A$, the probability that the DM chooses a in S, conditional on memory M. Randomness in choice may arise from two or more sources. First, retrieval from memory is probabilistic. For example, the likelihood that a particular experience can be retrieved from memory may depend on how often or how recently the corresponding alternative has been chosen in the past. Since retrieval is not deterministic, choice behavior that depends on memory may be non-deterministic as well. Second and irrespective of memory, value assessments in complex choice situations may be noisy, which too can lead to randomness in choice (Woodford 2020; Salant and Spenkuch 2024).

We assume that C satisfies two properties with respect to the value function V. First, conditional on memory, C is monotone in V. That is, for any two alternatives a and b with the same memory status, if $V(a, E) > V(b, E)$ then $C(S, a \mid M) \geq C(S, b \mid M)$; or, in words, the DM is responsive to differences in objective value. Second, C is continuous in V, so that, all else equal, small changes in objective value lead to small changes in choice probabilities.

[3] One may want to enrich the definition of an experience by including additional variables in it. Since we do not use additional variables in the analysis, we keep the definition as concise as possible.

To formalize how memory might affect choice behavior, let

$$\Delta(S, x \mid M, a) = C\left(S, x \mid M\right) - C\left(S, x \mid M^{-a}\right)$$

denote the change in the choice probability of alternative x in situation S if $a \in M$ was not in-memory. With this notation in hand, we define the memory premium as follows.

DEFINITION 1 (Memory Premium): *A choice function C gives rise to a* memory premium *if for any choice situation $S = (A, E)$, any memory M, and any alternative $a \in M$:*
(i) $\Delta(S, a \mid M, a) > 0$, and
(ii) $\Delta(S, a \mid M, a) > \Delta(S, b \mid M, a)$ for any alternative $b \neq a$.

In words, choice behavior exhibits a memory premium if an alternative is more likely to be chosen on account of being in- rather than out-of-memory; and if the increase in choice probability exceeds any potential increase in the choice probability of any other alternative.[4]

When C satisfies monotonicity and continuity with respect to V, Definition 1 implies that the DM is more likely to choose an in-memory alternative than an out-of-memory one, even when the objective value of the former is somewhat smaller than that of the latter. In particular, an in-memory alternative is more likely to be chosen than an equally good out-of memory alternative.

OBSERVATION 1: *Suppose C satisfies Definition 1, and is monotone and continuous in V. For any choice situation $S = (A, E)$, there exists a threshold $\epsilon_S > 0$ such that if $V(b, E) - V(a, E) \leq \epsilon_S$ for two alternatives $a \in M$ and $b \notin M$, then $C\left(S, a \mid M\right) > C\left(S, b \mid M\right)$.*

Several natural choice procedures can generate a memory premium. One example is a risk-neutral utility maximizer who perceives the value of in-memory alternatives to be systematically higher than it actually is. Another example is a utility maximizer who is risk-averse and holds more precise beliefs about in-memory alternatives compared to out-of-memory ones. A third example is a decision maker who maximizes from consideration sets that disproportionately include in-memory alternatives. A fourth example involves satisficing. A satisficer does not consider all possible alternatives to select the best one; instead, she sequentially examines a small number of alternatives and chooses the first one she finds satisfactory (Simon 1955). Such a DM would also display a memory premium if the order of consideration prioritizes in-memory alternatives. Of course, a hybrid choice procedure that combines satisficing and maximization may generate a memory premium as well. For

[4]Although the first part of Definition 1 is implied by the second one, we state part (i) for clarity.

instance, a DM who first evaluates whether some in-memory alternative is satisfactory before maximizing over the remaining options would similarly exhibit a memory premium.

2.3. *Comparative Statics*

The availability heuristic suggests that individuals judge the likelihood or frequency of events by the ease with which relevant instances come to mind (Tversky and Kahneman 1973). Extending this idea to choice behavior, we conjecture that the memory premium may also vary with the ease of recall. There are at least two reasons for why this may be the case. First, an alternative that is easier to retrieve from memory may be considered and evaluated prior to alternatives that are harder to recall or that are out-of-memory. When the DM satisfices, priority in consideration translates to a higher premium. Second, ease of recall may be interpreted by the DM as reflecting relevance, in which case the DM may assign higher value to alternatives that are easier to recall, and thus choose them more often. Drawing on previous findings in cognitive psychology and neuroscience, we now develop four testable implications that relate the size of the memory premium to well-established, fundamental regularities in the retrieval process.

First, memory decays over time. That is, recent experiences are more easily retrieved than temporally distant ones. If the memory premium is indeed linked to the ease or probability of recall, then we would expect a larger premium when the experience in question entered the DM's memory more recently.

PROPERTY 1 (Recency): *Consider two memories M_1 and M_2, which were obtained from memory M by adding to it a single experience $a \notin M$ in positions k_1 and $k_2 > k_1$, respectively. Then, for any choice situation S, $\Delta(S, a \mid M_1, a) > \Delta(S, a \mid M_2, a)$.*

A second fundamental property of memory is that repetition strengthens recall. Since more frequent experiences are more likely to be retrievable than less frequent ones, we would expect that repetition translates into a larger memory premium.

PROPERTY 2 (Repetition): *Let M_1 be a memory obtained from M by adding experiences with the alternative $a \notin M$ in positions $k_1 < \ldots < k_i$. Let M_2 be a memory obtained from M_1 by adding another experience with a in position $k_{i+1} > k_i$. Then, for any choice situation S, $\Delta(S, a \mid M_2, a) > \Delta(S, a \mid M_1, a)$.*

A third key property of memory is associativity. Past experiences are more readily retrieved when the context in which they occurred shares meaningful similarities with the current one. If choice behavior is, indeed, a function of recall, then the memory premium should be

larger when the current choice environment resembles the setting in which the in-memory alternative was previously chosen.

PROPERTY 3 (Similarity): *Consider two choice situations $S_1 = (A, E_1)$ and $S_2 = (A, E_2)$. If E_1 is more similar to the familiar environment than E_2, i.e., $\sigma(E_1, \hat{E}) > \sigma(E_2, \hat{E})$, then $\Delta(S_1, a \mid M, a) > \Delta(S_2, a \mid M, a)$ for any $a \in M$.*

If the memory premium is, indeed, linked to the ease with which an experience comes to mind, then similarity should have another measurable effect. Recent experiences stand out more when the associated environment is dissimilar to the choice environment in the past. Since salient experiences interfere with the recall of less salient ones, we should observe a smaller memory premium in subsequent choices when the distinctiveness of the most recent experience increases.

PROPERTY 4 (Interference): *Consider two unfamiliar environments, E_1 and E_2, such that E_1 is less similar to the familiar environment than E_2. Suppose the DM chose alternative a_i in environment E_i, and let M_i be the memory that is obtained as a result of attaching a_i to the first position of memory M.*[5] *If $x \in M$ then $\Delta(S, x \mid M_2, x) > \Delta(S, x \mid M_1, x)$.*

In words, Property 4 states that the memory premium ought to be smaller when the most recent choice environment stands out as different and thereby disrupts the recall of memories from the familiar environment.

In an analogous fashion to Property 4, imagine the following two scenarios. In the first, the DM encounters two unfamiliar environments, one immediately after the other. In the second scenario, the DM encounters the familiar environment $\hat{E}$ in between the two unfamiliar ones. We would expect that encountering the familiar environment has two effects. It should strengthen the DM's recall of other, prior experiences from the familiar environment, and it should interfere with her ability to recall the unfamiliar choice situation. As a consequence, we would expect a larger memory premium in the second scenario.

Our empirical analysis investigates whether choice behavior in unfamiliar environments does, indeed, exhibit a memory premium and, importantly, whether Properties 1–4 above hold. The context in which we study these questions is that of Chess960.

3. Setting and Data

3.1. *Chess960*

Chess960—also known as Fischer Random Chess—is a variant of standard chess, which was invented in 1996 by former world champion and legendary grandmaster Bobby Fischer.

[5] In symbols, $M_i = (a_i, a(1), b(2), \ldots)$, where $M = (a(1), b(2), \ldots)$.

The game uses the same board, pieces, and basic rules as regular chess. The key difference relative to standard chess is that the initial positions of the pieces on players' home ranks are randomized. That is, pawns are placed in their usual starting positions, but the pieces behind them are shuffled according to two rules: (i) Bishops must be placed on opposite-color squares, and (ii) the king must be positioned between the rooks to enable castling.[6] These constraints result in 960 possible initial board configurations—hence the name Chess960.

Chess960 provides an almost ideal environment to study the role of memory for decision-making in unfamiliar environments. First, there is no uncertainty about the rules of the game; yet first-time players of Chess960 may be unsure about the value of individual alternatives. After all, even minor differences in the initial placement of pieces can have important implications for optimal strategy. Second, individuals' memory can be empirically approximated. Chess players often memorize sequences of effective opening moves and counter-moves, on which they rely in standard chess. To the extent that we observe players' choices in standard chess, we can approximate which opening moves they have stored in memory. This, in turn, allows us to relate memory to their choices in the first game of Chess960. Third, randomness in the positioning of other pieces on the board provides us with exogenous variation in the quality of a given opening move. This variation is helpful in ruling out that in-memory alternatives are chosen more frequently because they happen to be better than out-of-memory ones. Fourth, data on games of chess and Chess960 are abundant, affording us enough statistical power to document even subtle relationships and effects.

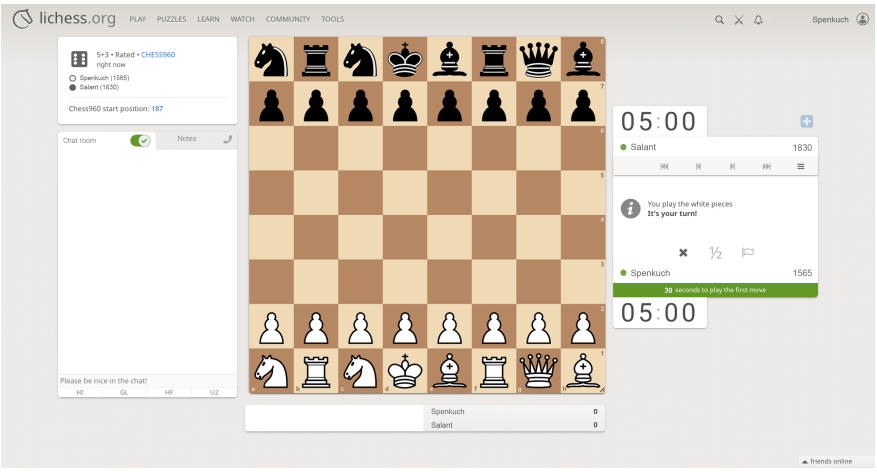
3.2. *Data and Descriptive Statistics*

Our data come from `lichess.org`, one of the most popular online chess servers.[7] Funded by donations, Lichess is ad-free and allows anyone to play live chess games at no cost through a high-quality graphical user interface. Every day, Lichess hosts several million games of standard chess and thousands of games of Chess960. For each variant, Lichess distinguishes between rated and casual games, both of which can be played under different time controls. Unlike causal games, rated games are consequential in the sense that their outcomes directly affect users' strength ratings and rankings on the site. They are thus only available to registered users. Since high ratings tend to be a source of pride among chess players, Lichess has a strict policy against computer-assisted play. Enforcement of this policy relies on a variety of methods, including community reporting of suspected offenders and automatic detection algorithms. Figure 1 shows a typical game of Chess960 on Lichess.

[6] In Chess960, the position of the king and rook after castling is exactly the same as in standard chess. Unlike in standard chess, this can mean that the king or rook "jump over" other pieces in the process of castling.

[7] Our description of these data borrows from Salant and Spenkuch (2024).

Figure 1: Chess960 on Lichess



Notes: Figure shows the beginning of a Chess960 game between registered users on `lichess.org`.

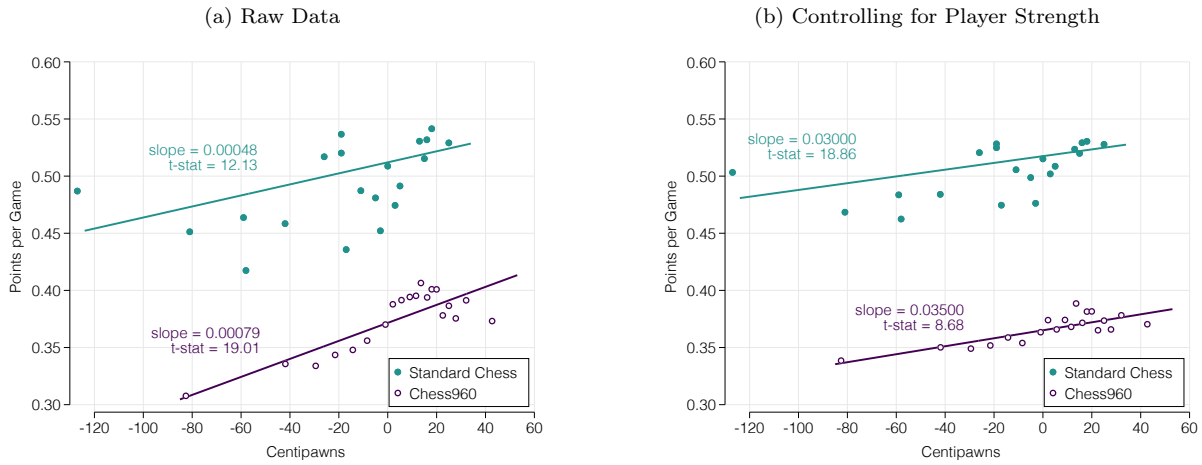
We have data on the universe of games on Lichess from January 2013 through June 2021. The available information includes players' usernames, strength ratings and real-world titles (if any), the type, date and start time of each game, the sequence of moves, as well as the ultimate outcome. For games of Chess960, we also observe the exact starting position. Given the granularity of our data, we can reconstruct every registered user's history of play on the platform. We, therefore, know every opening move of every player in standard chess games as well as in Chess960.

In our main analysis, we restrict attention to the opening move of approximately 147,000 players as they engage in a game of Chess960 for the first time.[8] There are two restrictions that limit the size of this sample. First, users must play their first game of Chess960 as White, and the game must be rated. While color is randomly assigned (in rated games), about half of first-time players choose to engage in a casual game. We nonetheless impose this restriction because we want to analyze choices with real stakes. Second, prior to their first game of Chess960, users must have executed at least twenty opening moves in standard chess games on Lichess. This restriction is meant to ensure that we can approximate which moves are stored in memory. In ancillary analyses, we study opening moves in players' first fifty games of Chess960 as White, irrespective of which color they played in the first game and whether that game was rated. We also impose higher thresholds on the number of previous standard games. Reassuringly, alternative sample restrictions yield qualitatively similar results.

Measuring the Quality of Opening Moves We complement the Lichess data with

[8]Appendix Figure AF.1 verifies that the distribution of starting positions in these games is consistent with random assignment.

Figure 2: Opening Moves and Game Outcomes



Notes: Figure shows binscatter plots of the relationship between game outcomes and the Stockfish ratings of opening moves in both standard chess and Chess960. Panel (a) does so based on the raw data, whereas panel (b) presents estimates of the same relationship after controlling for a fixed effect for a player's strength rating in standard chess interacted with that of her opponent. Wins are valued at one point, draws at half a point, and losses at zero.

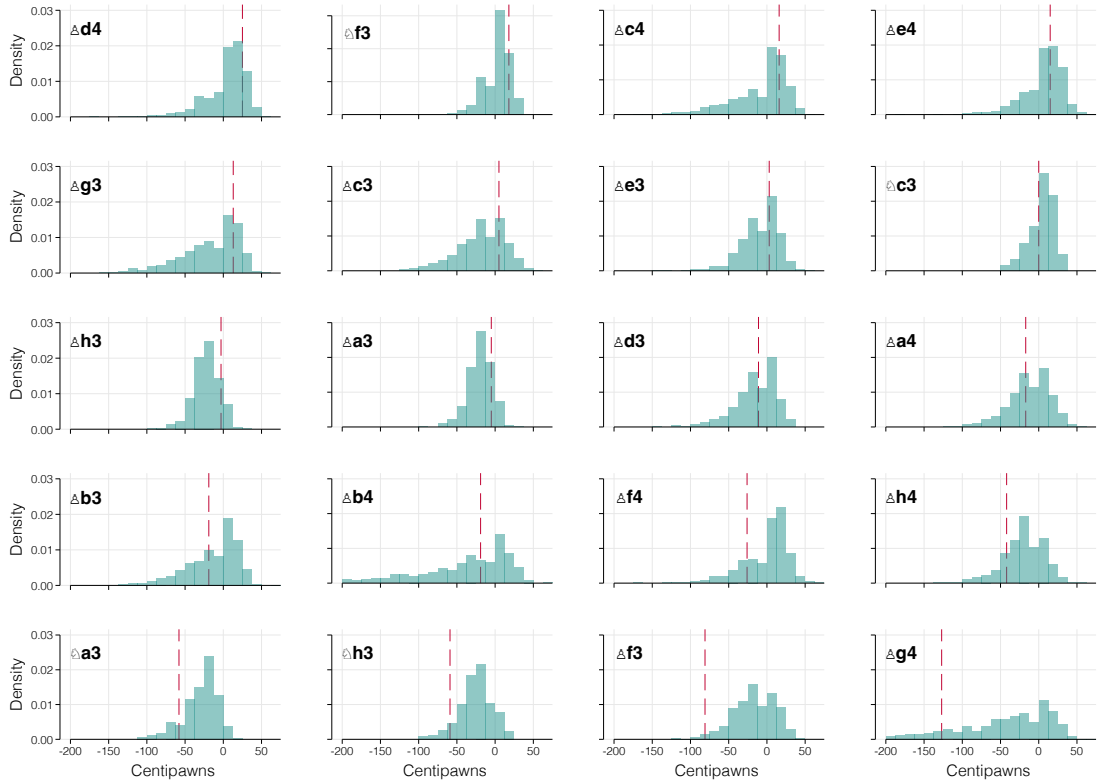
information on the objective quality of opening moves. To this end, we use the Stockfish software to rate every feasible opening move in every possible starting position. Stockfish is one of the most powerful open-source chess engines. It is far better at analyzing chess positions than even the best grandmasters.[9] Stockfish evaluates board configurations by examining sequences of possible moves (up to a certain depth) using a combination of heuristic evaluation functions and advanced neural networks. It then assigns a numerical score to each available move. A positive score reflects an advantage for White, whereas a negative score reflects an advantage for Black. The size of the advantage is measured in "centipawns," which correspond to one-hundredth of a pawn. For example, a move with a value of +50 leads to a position in which White is ahead by the equivalent of half a pawn.[10] Important for our purposes, Stockfish explicitly supports Chess960.

Figure 2 relies on Stockfish ratings to compare the importance of opening moves in standard chess and in the first game of Chess960. The left panel plots the relationship between the centipawn score of moves and the average outcome of the respective games in the raw Lichess data. Wins are valued at one point, draws at half a point, and losses are worth zero. The panel on the right shows the same relationship after carefully controlling for the strength of both players. In either panel, two patterns stand out. First, novice players of Chess960 win less than half of their games—at least in part because the new environment is unfamiliar and

[9]Recent versions of Stockfish can be expected to lose less than one percent of games against any human player.

[10]For comparison, commonly used rules of thumb value each bishop and knight at 300 centipawns, each rook at 500, and the queen at 900 centipawns.

Figure 3: Variation in the Quality of Opening Moves Across Starting Positions



Notes: Figure shows the distributions of centipawn scores of each opening move across all possible starting positions in Chess960. Panels are arranged according to the respective moves' ratings in standard chess, as indicated by the dashed vertical lines in each panel. Only opening moves that are available in standard chess are shown.

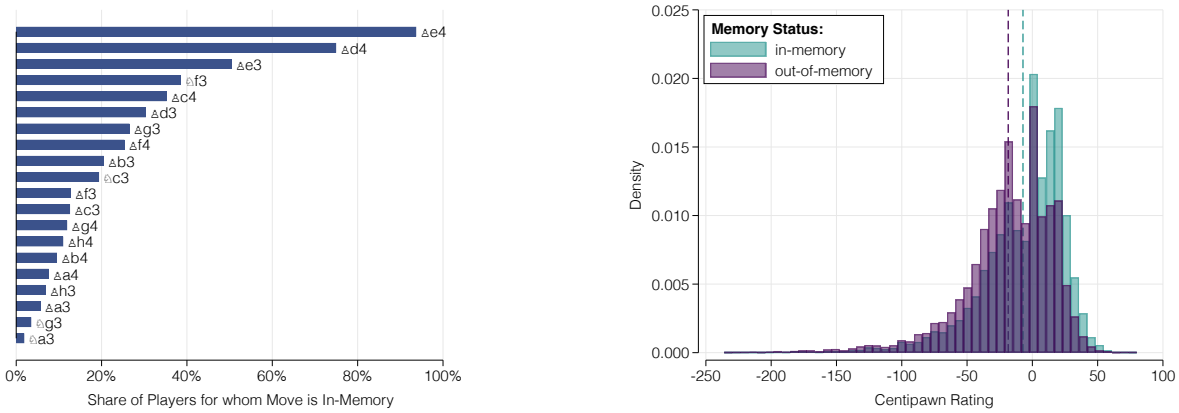
they lack the requisite experience. Second, and more importantly, there is a clear positive relationship between game outcomes and the quality of opening moves. This relationship is at least as strong in Chess960 as in regular chess. Taking the estimated slope in the right panel at face value, choosing an opening move that is approximately one standard deviation (≈ 35 centipawns) better increases White's chance of winning by about 1.2 percentage points.

Figure 3 shows that there is a great amount of within-move variation in quality across different starting positions. Each panel in this figure depicts the distribution of Stockfish ratings for a given opening move across starting positions in which the respective move is available. Vertical lines indicate an opening move's centipawn score in standard chess. Evidently, there are many starting positions in which popular opening moves, such as ♙d4 or ♙e4, are rated favorably; but there are also a number of positions in which playing one of these moves would put White at a significant disadvantage. Similarly, opening moves that tend to be bad in standard chess, such as ♙f3 or ♙g4, can, in fact, be quite good in Chess960. Overall, opening moves' centipawn scores in the standard starting position explain less than 6% of their variation in Chess960.

Figure 4: Memory

(a) Memory Status of Moves

(b) Quality of In- vs. Out-of-Memory Moves



Notes: Panel (a) shows the share of players for whom a particular opening move is in-memory, i.e., who have played this opening move at least once prior to their first game of Chess960. Panel (b) presents the distribution of centipawn ratings in players' first game of Chess960 for moves that are in- and out-of-memory. Vertical lines correspond to the mean rating in each category.

Measuring Memory To study the connection between memory and choice behavior, we must empirically approximate the content of players' memory, i.e., which moves they are able to recall. According to our workhorse definition, a given opening move is said to be "in-memory" for a particular player if she previously chose that move in at least one standard chess game on Lichess. In other words, we model memory based on individuals' entire history of play prior to their first game of Chess960, and we assume that they are able to recall any previous opening move.

An alternative approach would be to model recall as a function of the frequency and recency of play—and perhaps even outcomes. For instance, one might assume that memory is limited to the last one hundred games or to actions that were taken, say, at least twice. The rationale for such alternative definitions is that memory decays and that repetition strengthens recall. A significant drawback of the alternative approach is that it is somewhat arbitrary.

Since we do not know how long an experience remains in memory before it is forgotten, our workhorse definition labels any opening move that might, in principle, be recalled as in-memory. This is a conservative approach to estimating the memory premium. It is conservative because it will falsely classify some out-of-memory moves as in-memory, which will lead us to underestimate the true memory premium. As a robustness check, we estimate the memory premium for several thousand alternative definitions. Virtually all of them yield larger estimates (see Section 4).

Drawing on our workhorse definition, the left panel in Figure 4 shows the share of players for whom a particular opening move is in-memory. Unsurprisingly, ♙e4 and ♙d4—the most

Table 1: Summary Statistics

Variable	Mean	SD	Percentile 5%	25%	50%	75%	95%	*N*
A. Player Level								
ELO Rating in Standard Chess	1,555	344	1,030	1,300	1,530	1,790	2,150	147,357
Previous Chess Games	1,030	2,529	48	101	276	889	4,395	147,357
Previous Chess Games as White	521	1,284	24	51	139	448	2,224	147,357
In-Memory Moves	4.51	3.09	1	2	4	6	11	147,357
B. Game Level								
White Wins (%)	35.64	47.89						147,357
Draw (%)	3.31	17.88						147,357
White Loses (%)	61.04	48.77						147,357
Board Similarity	2.10	1.34	0	1	2	3	4	147,357
C. Move Level								
All Moves:								
Choose Move (%)	5.08	21.97						2,898,455
Stockfish Rating	–15.89	35.20	–80	–32	–11	9	27	2,898,455
In-Memory (%)	22.92	42.03						2,898,455
In-Memory Moves:								
Times Played in Chess	111.36	552.94	1	1	5	35	476	664,424
Days Since Last Played	71.08	168.10	0.01	0.94	9.33	56.13	357.78	664,424
Average Past Payoff	0.44	0.32	0.00	0.18	0.49	0.61	1.00	664,424

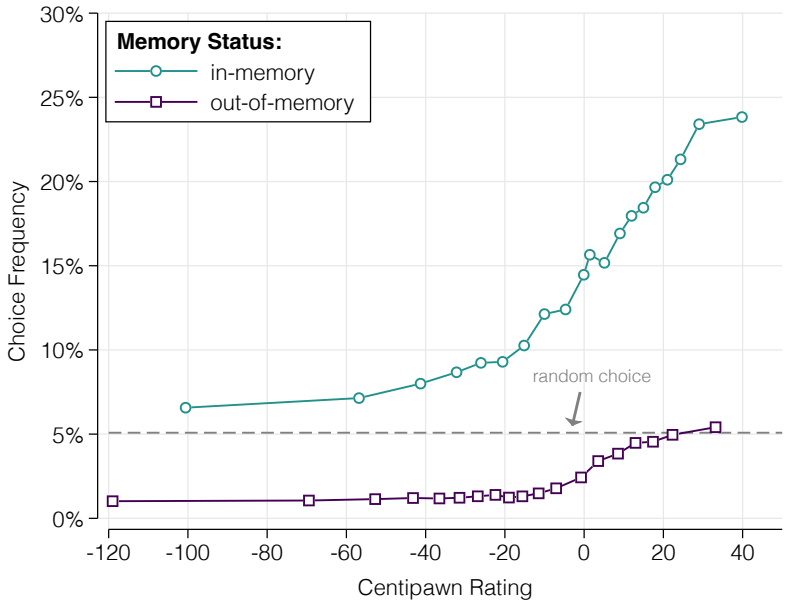
Notes: Table displays summary statistics for selected variables in our main sample. To be included in this sample, players must play their first game of Chess960 as White, and the game must be rated. We further require that they play at least twenty games of standard chess on Lichess prior to their first game of Chess960. Observations in Panel A correspond to players. Each observation in Panel B corresponds to a player's first game of Chess960. Panel C reports move-level information, so that each observation corresponds to a feasible opening move in a particular game.

common opening moves in standard chess—are in-memory for the vast majority of players.[11] By contrast, a total of eleven other moves are out-of-memory for more than four out of five individuals. The right panel in Figure 4 presents the distribution of Stockfish ratings across starting positions for players' in- and out-of-memory moves. Although both distributions overlap significantly, in-memory moves have, on average, higher ratings.

Summary Statistics Descriptive statistics for the most important remaining variables in our data are shown in Table 1. The average (median) player has engaged in more than 1,000 (270) standard chess games prior to her first game of Chess960. Over the course of these games, she has executed 4.5 (4) different opening moves, which she now holds in memory. The individuals in our analysis thus have substantial experience playing chess. Yet 61% of them lose their first game of Chess960. The latter observation suggests that Chess960 is, indeed, unfamiliar.

[11] Opening moves that are only available in Chess960 are never classified as in-memory. They are excluded from Figure 4. To gauge the popularity of each opening move in standard chess, see Appendix Figure AF.2.

Figure 5: Choice Frequencies in the Raw Data, by Memory Status



Notes: Figure shows binscatter plots of the relationship between the centipawn ratings and choice frequencies of in- and out-of-memory moves in players' first game of Chess960 as White.

4. The Memory Premium

Figure 5 provides the first piece of evidence that in-memory alternatives are chosen more frequently than equally good out-of-memory ones. The figure depicts binscatter plots relating move quality to choice frequency. It does so for in-memory and out-of-memory moves. Two key patterns emerge. First, conditional on memory status, choice frequencies are monotonically increasing in quality, as assumed in Section 2. Second, comparing moves with similar centipawn ratings, the choice frequency of in-memory alternatives exceeds that of out-of-memory ones by anywhere from five to twenty percentage points. Strikingly, the worst in-memory moves are chosen at a slightly higher rate than the best out-of-memory ones, suggesting that ϵ_S in the statement of Observation 1 is meaningfully large.

To provide more rigorous evidence of the memory premium, we estimate the following linear probability model:

$$(1) \qquad Choose_{p,m,s} = \delta\, In\text{-}Memory_{p,m} + \gamma\, Quality_{m,s} + \mu_m + \epsilon_{p,m,s}.$$

Here, $Choose_{p,m,s}$ is an indicator for whether player p selected opening move m in starting position s, $In\text{-}Memory_{p,m}$ denotes whether p holds m in memory, $Quality_{m,s}$ corresponds to the move's objective value, as measured by its Stockfish rating in the relevant starting position, and μ_m is an opening-move fixed effect. The latter accounts for the possibility

Table 2: Memory Premium

	Probability of Choosing Move					
	(1)	(2)	(3)	(4)	(5)	(6)
In-Memory (÷ 100)	12.07	3.83	3.83	4.52	4.52	4.53
	(0.04)	(0.04)	(0.04)	(0.05)	(0.05)	(0.05)
Centipawn Rating (÷ 100)		0.03				
		(0.00)				
Fixed Effects:						
Move	No	Yes	No	No	No	No
Move × Position	No	No	Yes	No	No	No
Move × Position × Size of Memory	No	No	No	Yes	Yes	Yes
Player × Opponent Strength	No	No	No	No	Yes	Yes
Previous Standard Games as White	No	No	No	No	No	Yes
Mean of LHS Variable	5.08	5.08	5.08	5.08	5.08	5.08
R^2	0.05	0.13	0.16	0.24	0.24	0.24
N	2,898,455	2,898,455	2,898,455	2,898,455	2,898,455	2,898,455

Notes: Entries are coefficients and standard errors from estimating variants of the regression model in eq. (1) by ordinary least squares. The sample consists of all available opening moves in players' first game of Chess960 as White. All estimates are scaled to correspond to the percentage-point change in choice probability associated with a one-unit increase in the respective regressor. Standard errors are clustered by player, and are shown in parentheses.

that some moves may be more likely to be chosen because they accord with popular advice like "open with a center pawn." The coefficient of interest is δ. It is identified by comparing the same opening move across players for whom a particular opening move is and is not in memory, holding the objective quality of the move as well as its latent popularity fixed.

In more granular specifications, we interact the move fixed effect with indicators for the starting position and the size of players' memory, i.e., the number of other in-memory moves. In these specifications, δ is identified by comparing the same opening move in the same starting position across players that hold the same number of other moves in memory. These comparisons fix the quality of the respective move, the set of available alternatives, as well as how many other opening moves a player may recall.[12] They need not, however, involve players that are otherwise similar. In order to address potential concerns about systematic differences across players, we estimate models that also control for player characteristics, such as their prior chess experience and skill. In Section 6, we additionally leverage within-player variation over time. For the remainder of this section, however, we focus on players' first move in their first game of Chess960.

Table 2 presents our main result. The coefficient in col. (1) indicates that, in the raw data, the average choice frequency of in-memory moves exceeds that of out-of-memory ones by 12.1 p.p. Controlling for the objective quality as well as the latent popularity of a move narrows

[12]In Appendix Table AT.1, we replicate our main result controlling for *which* other moves players hold in memory. The implied comparisons come even closer to approximating the definition of the memory premium in Section 2, but yield point estimates and comparative statics that are less precise. Reassuringly, both sets of results are qualitatively equivalent.

the difference between both types of alternatives to about 3.8 p.p. (col. 2). The point estimate shrinks because opening moves that are generally more likely to be played—e.g., because they help to gain control of the center of the board—are also more likely to be in memory. The next column also accounts for the set of available alternatives by including move-by-starting-position fixed effects. Holding the choice set as well as the quality of all available alternatives fixed has no appreciable effect on the coefficient of interest. The specification in col. (4) additionally controls for the size of players' memory, which increases the estimated memory premium to about 4.5 p.p. Columns (5) and (6) add player-level controls, i.e., players' strength ratings in regular chess (interacted with those of their opponents) and their prior experience. Accounting for experience is potentially important because, given our definition of memory, players with longer histories of play tend to have different memories than those with shorter histories. Reassuringly, controlling for systematic differences across players has almost no effect on the estimated coefficients.

In sum, the results in Table 2 document that a given opening move is about 3.8–4.5 p.p. more likely to be chosen when it is in- rather than out-of-memory. Given a mean choice frequency of about 5%, our findings imply a memory premium that is economically very large.

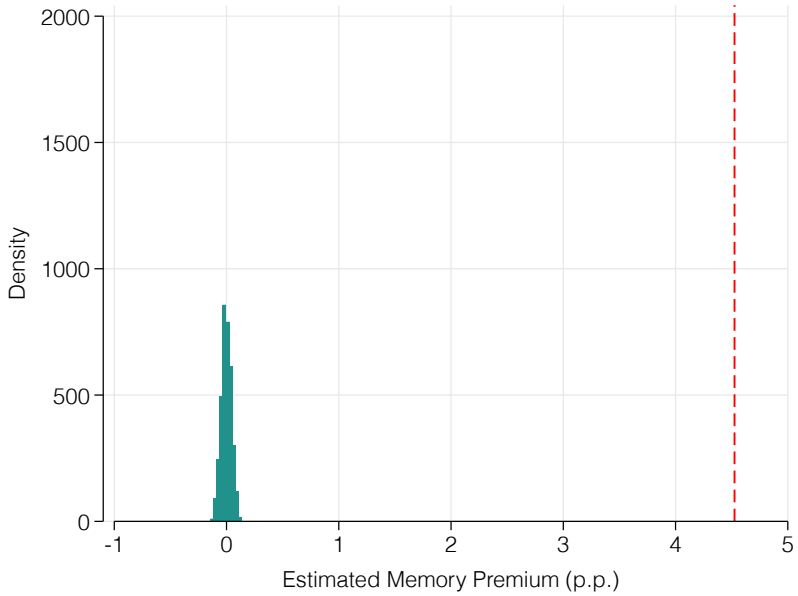
Robustness To assess the robustness of our main result we conduct a placebo test. Specifically, we randomly reassign histories of prior play across the individuals in our data. We then re-estimate the regression model in col. (6) of Table 2 based on the newly induced memories. Repeating this process a thousand times, we compare the true point estimate to the distribution of placebo coefficients.

Figure 6 shows the results. As one might expect, the median placebo estimate is approximately zero.[13] The remaining placebo coefficients are tightly clustered around the median. As a result, the true point estimate from col. (6) of Table 2 is a clear outlier. It is more than twenty times larger than the largest placebo coefficient, which suggests that players' memory and choice behavior are, indeed, linked.

We proceed to assess the robustness of our main finding to different definitions of memory. To this end, we follow the alternative approach outlined in Section 3. We say that an opening move is in-memory if, in the last n games, a player has used it at least t times. Each alternative definition is, therefore, characterized by a pair (n, t), where n denotes how many games players can hold in memory (i.e., its length) and t is a usage threshold that governs which moves can be recalled. Using this notation, our workhorse definition of memory corresponds

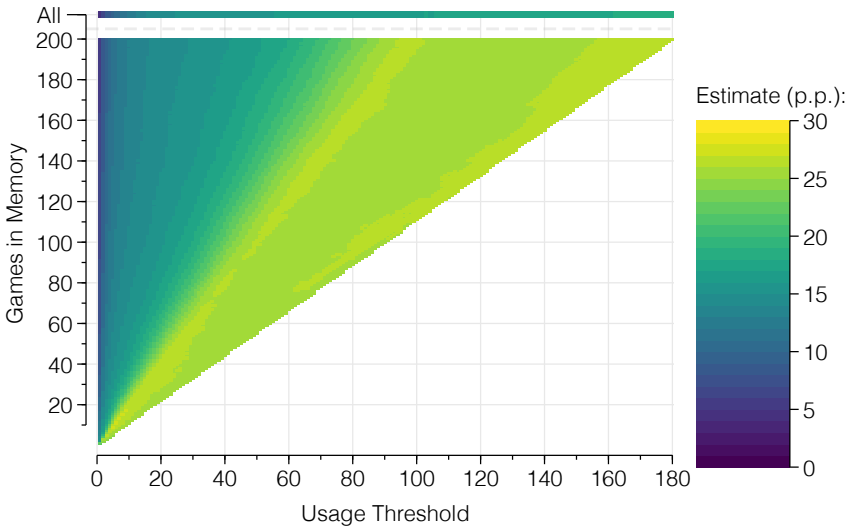
[13]By contrast, the regression model in col. (1) of Table 2 produces placebo estimates that tend to be significantly greater than zero. This observation is consistent with omitted variables bias, which explains why the point estimate decreases upon controlling for centipawn ratings and move fixed effects.

Figure 6: Distribution of Placebo Estimates



Notes: Figure shows the distribution of placebo estimates of the memory premium. Placebo estimates are derived from randomly reshuffling the observed memories across players and estimating the regression model in col. (6) of Table 2 on the same sample of moves. The vertical line corresponds to the original estimate of the memory premium in that column.

Figure 7: Estimated Memory Premia under Alternative Definitions of Memory



Notes: Figure shows estimates of the memory premium under alternative definitions of memory. Each alternative definition corresponds to a pair (n, t), for which a given move is said to be in-memory if it has been played at least t times (x-axis) in the last n games of standard chess (y-axis). All estimates are based on the regression model and sample in col. (6) of Table 2, excluding players who have not played sufficiently many prior games for any moves to be classified as in-memory.

Table 3: Heterogeneity in the Memory Premium

	Probability of Choosing Move					
	(1)	(2)	(3)	(4)	(5)	(6)
In-Memory (÷ 100)	4.66	4.41	3.15	6.49	6.41	3.57
	(0.07)	(0.07)	(0.06)	(0.09)	(0.10)	(0.05)
Sample	Strong Players	Weak Players	High Experience Players	Low Experience Players	Close to Optimal Moves	Suboptimal Moves
Fixed Effects:						
Move × Position × Size of Memory	Yes	Yes	Yes	Yes	Yes	Yes
Player × Opponent Strength	Yes	Yes	Yes	Yes	Yes	Yes
Previous Standard Games as White	Yes	Yes	Yes	Yes	Yes	Yes
Mean of LHS Variable	5.08	5.08	5.08	5.08	9.35	3.35
R^2	0.30	0.32	0.31	0.28	0.25	0.23
N	1,456,160	1,442,295	1,446,770	1,451,685	839,173	2,059,282

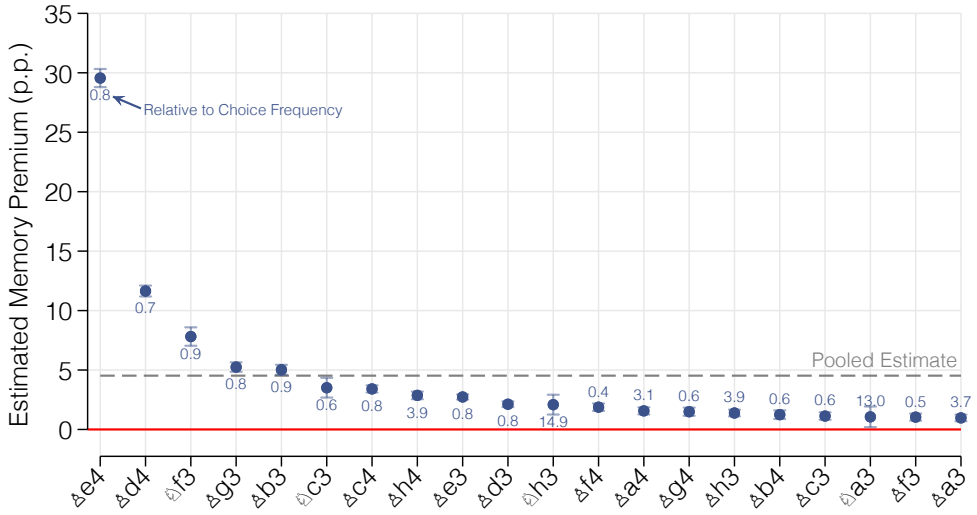
Notes: Entries are coefficients and standard errors from estimating the most saturated regression model in Table 2 by ordinary least squares. The sample consists of all available opening moves in players' first game of Chess960 as White. Columns 1 and 2 distinguish between players with above and below median strength ratings in standard chess. Columns 3 and 4 split the sample at the median of the number of previous standard chess games on Lichess, whereas columns 5 and 6 distinguish between moves' whose Stockfish rating does and does not fall within 20 centipawns of that of the best available alternative in the relevant starting position. All estimates are scaled to correspond to the percentage-point change in choice probability associated with a one-unit increase in the respective regressor. Standard errors are clustered by player, and are shown in parentheses.

to $n = \infty$ and $t = 1$. For each player in our data, we construct alternative memories for all possible combinations of $n \leq 200$ and $1 \leq t \leq 0.9n$, and for $n = \infty$ with various usage thresholds. We then estimate the memory premium under each alternative definition.

Figure 7 visualizes the resulting point estimates. Anticipating the comparative statics in the next section, we note two regularities. First, holding the length of players' assumed memory fixed, we tend to observe larger coefficients for higher usage thresholds, i.e., when in-memory moves have been used more frequently in the past. Second, for any particular usage threshold, we find that the estimated memory premium decreases if we assume that more distant openings can still be recalled—consistent with memory decay. More importantly, the estimates in Figure 7 range from about 4.5 to 27 p.p. Comparing this range with the coefficients in Table 2 demonstrates that our workhorse definition yields a conservative estimate of the memory premium.

Heterogeneity Returning to our workhorse definition, Table 3 investigates heterogeneity in the memory premium. The first two columns present estimates for players with above- and below-median strength ratings in regular chess. The next two columns differentiate players according to their previous chess experience. Finally, the last two columns distinguish between opening moves that are and are not close to optimal in a given starting position (i.e., within 20 centipawns of the best available move). While we do not detect meaningful heterogeneity by skill level, we do find that the memory premium is about twice as large for players who

Figure 8: Heterogeneity in the Memory Premium Across Opening Moves



Notes: Figure shows estimates of the memory premium for each opening move in standard chess. All estimates are based on the regression model and sample in col. (6) of Table 2, allowing for heterogeneity in the memory premium across opening moves. Error bars correspond to 95%-confidence intervals, accounting for clustering by player.

have played fewer than the median number of regular chess games on the platform. We also find heterogeneity by relative move quality. The estimated memory premium is significantly higher for moves that are close to optimal. Yet the point estimate in col. (6) of Table 3 indicates that, even for moves that are suboptimal, the memory premium is economically large and statistically different from zero.

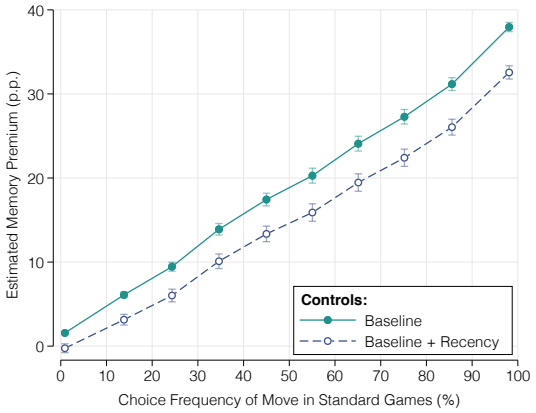
We also investigate heterogeneity across individual moves. Figure 8 plots the estimated memory premium for each of the twenty opening moves in standard chess. Although we observe much larger premia for moves that are frequently used in regular chess, such as ♙e4, ♙d4 and ♘f3, it is important to note that all estimates in Figure 8 are statistically significant and economically sizeable—especially compared to the mean choice frequencies of the respective moves. Appendix Table AT.2 further distinguishes between starting positions in which a particular move is and is not close to optimal. In the former case, 17 out of 20 point estimates are positive and statistically significant.[14] In the latter case, all twenty estimates are positive and statistically different from zero. The memory premium, therefore, applies to all moves, not just ones that are popular or good.[15]

[14]Of the remaining three estimates, two are positive and one is negative. All three are statistically insignificant.

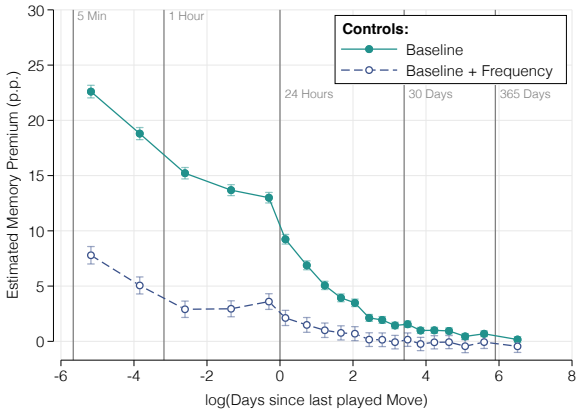
[15]In Appendix D, we address the concern that our findings might be driven by players attempting to execute particular opening strategies. Based on the evidence in Appendix Tables AT.3 and AT.4 , we conclude that such tendencies are not quantitatively important. We further note that any concern about the memory premium that is specific to chess cannot explain the experimental results in Section 7.

Figure 9: Reinforcement and Decay

(a) Memory Premium, by Frequency of Prior Use



(b) Memory Premium, by Recency of Prior Use



Notes: Figure presents estimates of the memory premium by frequency (left panel) and recency (right panel) of prior use of the move in standard chess. Solid lines and markers correspond to estimates based on the regression model and sample in col. (6) of Table 2. Dashed lines and hollow markers correspond to estimates that additionally control for recency (left panel) and frequency (right panel) of prior use. Error bars show 95%-confidence intervals, accounting for clustering by player.

5. Comparative Statics

In order to establish a more-direct connection between the memory premium and the ease with which experiences can be retrieved from memory, we proceed to examine Properties 1–4. That is, we ask whether the memory premium exhibits the same key comparative statics that cognitive scientists have previously documented for the probability of recall.

5.1. *Repetition and Decay*

A fundamental property of memory is that it is reinforced by repetition, and that it decays over time. Figure 7 above already provides some evidence that similar forces operate on the memory premium. After all, alternative definitions of memory that impose higher usage thresholds and limit recall to shorter histories of play yield larger point estimates. To investigate the comparative statics of the memory premium more carefully, we apply our workhorse definition and extend our most saturated regression model in Table 2 to allow for heterogeneity in δ, depending on how often and how recently a given in-memory move had previously been used.

Figure 9 presents the results. The solid line and markers in the panel on the left show estimates of the memory premium according to the frequency with which the respective player chose the move in regular chess. The dashed line and hollow markers correspond to estimated premia after additionally controlling for how recently the respective move had been played. The evidence in this panel indicates that the memory premium increases greatly with the frequency of prior use. Property 2 is, therefore, supported by the data.

The right panel of Figure 9 shows how the memory premium varies with recency, with and

without controlling for the frequency with which a player relied on the respective move in the past. Consistent with Property 1, moves that had been played more recently in standard chess exhibit significantly higher memory premia.[16]

In Appendix Figure AF.5, we differentiate between moves with above- and below-median payoffs, i.e., moves that a given player used more or less successfully in the past (as measured by the outcomes of the relevant games). Conditional on a move's frequency of use, the memory premium increases in prior success—though perhaps not as much as one might have expected. Although there is a strong unconditional relationship between choice frequencies in Chess960 and moves' prior payoffs, this relationship appears to be mediated by the frequency of use. Taken together, the evidence suggests that the memory premium is strengthened by repetition and that it decays over time.

5.2. *Associativity*

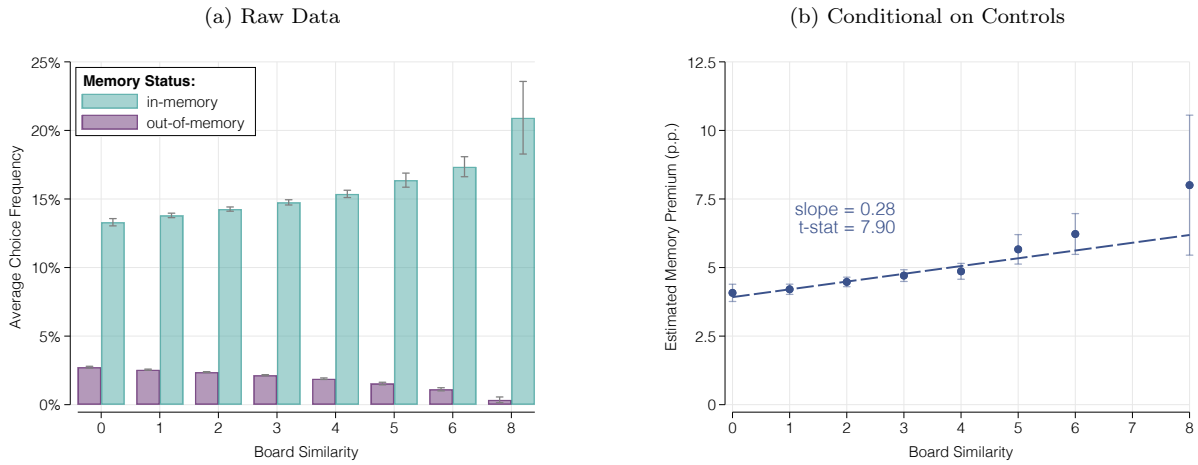
Memory is associative in the sense that past experiences are more likely to be recalled when they are related to current environmental cues. To investigate this property, we rely on the fact that starting positions in Chess960 are randomized. We define the similarity between a particular position in Chess960 and the standard chess board as the number of pieces that are initially placed on the same squares. If the memory premium is linked to the probability of recall, then we would expect to estimate higher premia in starting positions with greater resemblance to the standard one.

This is, indeed, what we find. The left panel of Figure 10 partitions the raw data according to board similarity, and shows raw choice frequencies for in- and out-of-memory moves.[17] The right panel shows, for each level of similarity, regression estimates of the memory premium. Even after controlling for players' prior experience and skill and accounting for the quality of all available alternatives via position-by-move-by-size-of-memory fixed effects, we observe larger memory premia in starting positions that resemble the standard chess board more. The evidence, therefore, suggests that the memory premium is associative, consistent with Property 3.

[16] Appendix Figure AF.4 replicates the analysis in the right panel of Figure 9 based on "donut definitions" of memory. According to these definitions, a move is said to be in-memory if, and only if, it had been played at least once more than, say, a month prior to a player's first game of Chess960. This robustness check helps to address concerns about short-term serial correlation in the choice of opening moves. For replications of our main result in Table 2 based on donut definitions of memory, see Appendix Tables AT.5–AT.8.

[17] Given that the standard chess board is one of the possible starting positions, our measure of similarity ranges from zero to eight. There are no starting positions in which seven pieces are placed on their usual squares.

Figure 10: Memory Premium, by Board Similarity



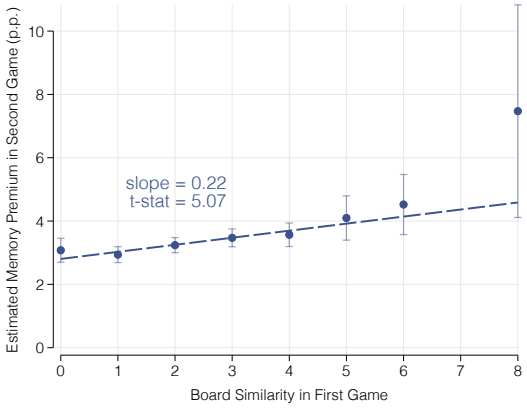
Notes: Panel (a) shows average choice frequencies of in- and out-of-memory moves by similarity between the starting position in Chess960 and that in standard chess. Similarity is defined as the number of pieces that are placed on the same squares as in standard chess. Panel (b) presents estimates of the memory premium by board similarity. These estimates are based on the regression model and sample in col. (6) of Table 2, allowing for heterogeneity in the memory premium by board similarity. Error bars correspond to 95%-confidence intervals, accounting for clustering by player.

5.3. *Interference*

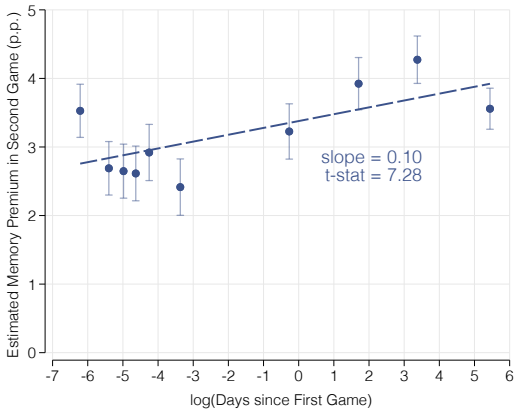
Experiences compete for recall to the point where new experiences interfere with the retrieval of older ones. To test for interference effects in the memory premium, we take a three-pronged approach. First, we test Property 4 by asking whether the memory premium in a player's *second* game of Chess960 as White depends on the similarity between the standard chess board and the starting position in the *first* game. We would expect that memories of the first game of Chess960 are more salient when the starting position was very different from the standard one. They should thus provide greater competition when it comes to recalling experiences from standard chess. If correct, then the memory premium in the second game ought to be an increasing function of the similarity between the standard starting position and that in the player's first game of Chess960. Second, we ask whether the memory premium in the second game of Chess960 as White increases in the temporal distance to the first one. A long delay after the first game will cause the associated memories to decay. They should, therefore, interfere less with the retrieval of other experiences. Third, we investigate whether the memory premium in the second game increases in the number of standard chess games between the first game of Chess960 and the second one, conditional on board similarity in the first game and the time that has elapsed since. The rationale for our third test is that newer experiences from standard chess will directly interfere with recalling the first game of Chess960 while strengthening the recall of experiences from previous standard games. This, in turn, should increase the memory premium.

Figure 11: Evidence of Interference

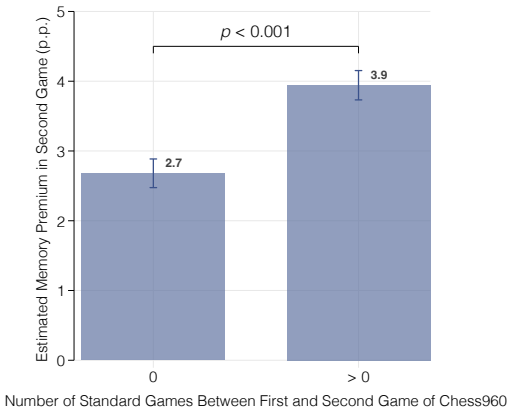
(a) Interference via Similarity



(b) Interference over Time



(c) Interference by Newer Memories



Notes: Panel (a) presents estimates of the memory premium in players' second game of Chess960, by similarity between the starting position in their first game of Chess960 and the starting position in standard chess. Panel (b) shows estimates of the memory premium in players' second game of Chess960, by decile of elapsed time between their first and second Chess960 games. Panel (c) distinguishes between the memory premium in the second game of Chess960 when players have and have not played any standard chess games between their first and second games. The estimates in panels (a) and (b) are based on the regression model in col. (6) of Table 2, allowing for heterogeneity in the memory premium along the relevant dimension. The estimates in panel (c) additionally control for board similarity in the first game of Chess960, as well as the time that elapsed between the first two games. The sample in each panel consists of all available opening moves in players' second game of Chess960, restricting attention to cases in which each of the first two games was rated and played as White. Error bars correspond to 95%-confidence intervals, accounting for clustering by player.

Figure 11 provides evidence in support of all three predictions. Measuring board similarity in the same way as before, panel (a) of Figure 11 shows that the memory premium in the second game is significantly smaller when the starting position in the first game of Chess960 was very different from the standard one. Panel (b) presents estimates of the memory premium in the second game as a function of the time between the first game and the second one. The positive slope indicates that the memory premium increases as more time elapses. Panel (c) compares estimates of the memory premium when players have and have not engaged in at least one game of standard chess between their first two games of Chess960.[18] The memory premium is significantly smaller when they have not. The evidence in panels (a)–(c) is, therefore, consistent with the idea that the memory premium is moderated by interference in recall.

[18]About 60% of individuals in our data play zero standard games between their first two games of Chess960.

6. Dynamics and Causality

We next examine how the memory premium evolves as individuals become more familiar with the new choice environment. We also provide real-world evidence to suggest that the connection between memory and choice behavior is plausibly causal.

6.1. *Dynamics of the Memory Premium*

In order to investigate how the memory premium evolves as decision-makers gain experience, we change the scope of our analysis to all individuals who played at least fifty games of Chess960 as White, regardless of which color they were assigned in their first game and whether that game was rated. There are 16,731 such players in our data. We then estimate the most saturated regression model in Table 2 separately for each of their first fifty games. In doing so, we distinguish between the memory premium in all games (incl. casual ones) and in rated games only. Drawing on all games affords us slightly more statistical power but comes with the drawback of lower stakes.

Figure 12 shows how the memory premium evolves over time. Irrespective of whether we consider all games or only rated ones, two patterns stand out. First, despite limited overlap in samples, the estimated memory premium in players' first game of Chess960 as White is quantitatively similar to that in Table 2. Second, the memory premium decreases rapidly as players gain experience; but it does not vanish. In fact, after about twenty-five games, the estimated memory premium stabilizes around 2.5 p.p., or about 50% of the average choice frequency in this setting.[19]

6.2. *Causality*

Does being retrievable from memory *cause* higher choice frequencies? In order to provide evidence in support of a causal relationship between memory and choice behavior, we consider a dynamic definition of memory. According to this definition, a move is said to be in-memory if, and only if, a player had chosen it at least once in a standard chess game on Lichess prior to the *current* game of Chess960. A move can thus enter a player's memory between her first game of Chess960 and subsequent ones.[20] Does the frequency with which the respective move is chosen in Chess960 increase afterwards?

[19]Appendix Figure AF.6 replicates the analysis in Figure 12 based on a definition of memory that is based on players' choices in prior games of Chess960. Although we estimate a significantly larger premium for this definition of memory, we observe a qualitatively similar decline as players gain experience. The observation that the memory premium is higher for moves that have previously been used in Chess960 is consistent with associative recall.

[20]By contrast, our (static) workhorse definition of memory considers only opening moves in standard games prior to the *first* game of Chess960 (see Section 3).

Drawing on the 41,990 instances in which an opening move enters a player's memory after her first game of Chess960, Figure 13 plots the average choice frequency of these moves before and after they were first used in a regular chess game. Consistent with the idea that players experiment more with out-of-memory alternatives as they gain experience, choice frequencies start out below average and increase over time. Importantly though, immediately after a move is used for the first time in a standard game, the frequency with which it is chosen in Chess960 jumps by approximately 0.3 p.p., or about 6%.[21]

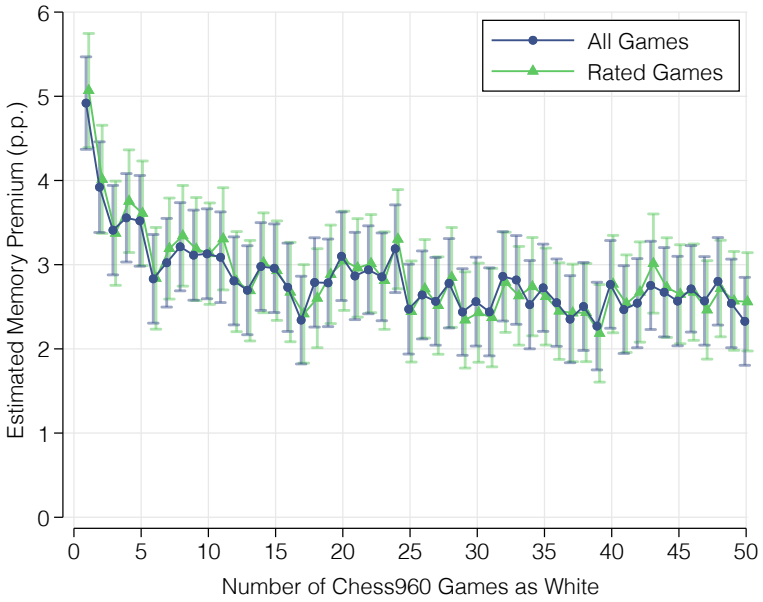
The upper panel of Table 4 shows that the discontinuity in Figure 13 is robust to different estimators, bandwidths, and samples. For instance, the bias-corrected RD estimator of Calonico et al. (2024), in conjunction with the corresponding optimal bandwidth, yields point estimates of 0.37 p.p. in the sample of all games and 0.36 p.p. when we restrict attention to rated games only. Relying on the standard linear model and alternative bandwidths produces estimates ranging from 0.22 to 0.29 p.p., all of which are statistically significant.

To be clear, these estimates average over situations in which a given opening move had truly never been used before—either in regular chess or Chess960—and those in which it had not been played in standard games on Lichess but was, in fact, used in games outside of the platform or in an individual's first few games of Chess960. In the latter cases, we should not think of the move as literally entering memory. Rather, it seems more appropriate to interpret the event of a move's first use in a standard game on Lichess as reinforcing prior experiences, which, in turn, improves recall. Either way, the observed discontinuity suggests that retrievability from memory affects choice behavior.

The lower panel of Table 4 complements our RD estimates with evidence from a difference-in-differences approach. The former research design identifies the effect of moves "entering memory" by comparing average choice frequencies *just* before and after the event. Some of the relevant variation, therefore, comes from comparisons across different players and moves. The latter design compares changes in the choice frequency of moves that enter players' memories to contemporaneous changes in the use of other moves. That is, the difference-in-differences approach leverages variation *within* "treated" and "untreated" player-move combinations. Reassuringly, both research designs yield qualitatively equivalent answers.

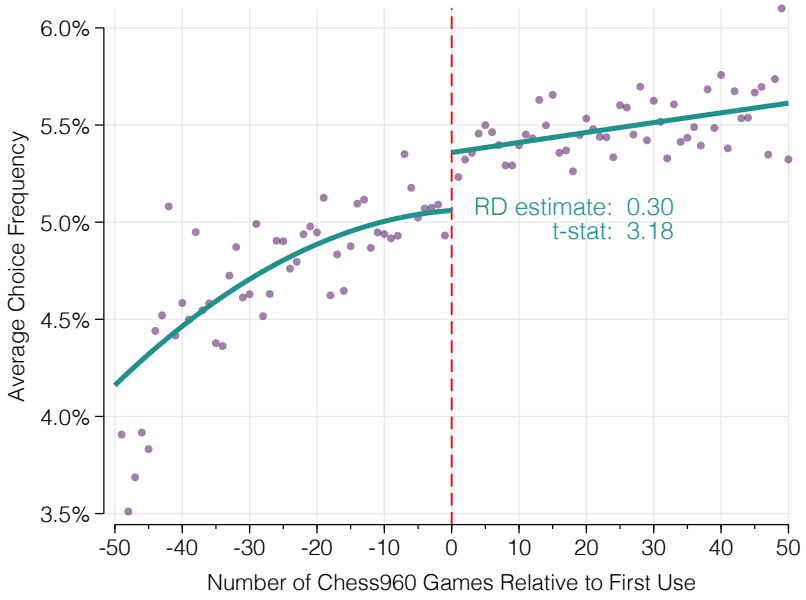
[21] Appendix Figure AF.7 shows the density of observations around the threshold. There is no apparent discontinuity.

Figure 12: Dynamics of the Memory Premium



Notes: Figure shows two series of estimates of the memory premium in each of players' first fifty games of Chess960 as White. The first series restricts attention to rated games, whereas the second series considers both rated and unrated games. All estimates are based on the regression model in col. (6) of Table 2. Error bars correspond to 95%-confidence intervals, accounting for clustering by player.

Figure 13: Evidence from an RD Design



Notes: Figure plots average choice frequencies of opening moves in Chess960 as a function of the number of Chess960 games as White relative to the respective move's first use in standard chess. Each bin averages over one game. Opening moves that had been used prior to a player's first game of Chess960 are excluded. Solid lines correspond to a second-order polynomial, fitted separately on each side of the threshold.

Table 4: Estimates from Causal Research Designs

A. RD Estimates

	Probability of Choosing Move					
	(1A)	(2A)	(3A)	(4A)	(5A)	(6A)
In-Memory (÷ 100)	0.25	0.22	0.29	0.26	0.37	0.36
	(0.09)	(0.10)	(0.11)	(0.12)	(0.14)	(0.16)
Sample	All Games	Rated Games	All Games	Rated Games	All Games	Rated Games
Estimator	OLS	OLS	OLS	OLS	Calonico et al. (2014)	Calonico et al. (2014)
Polynomial	Linear	Linear	Linear	Linear	Linear	Linear
Bandwidth	20	20	10	10	Optimal	Optimal
Mean of LHS Variable	5.26	5.28	5.24	5.27	5.25	5.29
N	1,219,263	1,080,203	687,013	602,314	565,062	493,419

B. Difference-in-Differences Estimates

	Probability of Choosing Move					
	(1B)	(2B)	(3B)	(4B)	(5B)	(6B)
In-Memory (÷ 100)	0.43	0.36	0.43	0.42	0.26	0.21
	(0.05)	(0.05)	(0.07)	(0.07)	(0.06)	(0.06)
Sample	All Games	Rated Games	All Games	Rated Games	All Games	Rated Games
Estimator	TWFE	TWFE	Borusyak et al. (2024)	Borusyak et al. (2024)	Cengiz et al. (2019)	Cengiz et al. (2019)
Fixed Effects:						
Player × Move	Yes	Yes	Yes	Yes	No	No
Chess960 Games as White	Yes	Yes	Yes	Yes	No	No
Player × Move × Stack	No	No	No	No	Yes	Yes
Chess960 Games as White × Stack	No	No	No	No	Yes	Yes
Mean of LHS Variable	5.08	5.08	3.47	3.45	3.27	3.27
N	16,453,129	14,704,058	12,137,434	10,734,051	15,501,782	13,833,500

Notes: Entries in Panel A are estimates of the memory premium from a regression discontinuity design around the first time a player uses a particular opening move in standard chess. Columns 1A–4A report estimates based on a standard linear regression model, restricting attention to observations within the respectively indicated bandwidths around the threshold and allowing for different slopes in the running variable on either side of the discontinuity. Columns 5A and 6A report bias-corrected regression discontinuity estimates based on the procedure of Calonico et al. (2014), using a uniform kernel and the associated MSE-optimal bandwidth. The sample in Panel A consists of all available opening moves in players' first fifty games of Chess960 as White, provided that they enter players' memory after the first game of Chess960. A move is said to enter a player's memory when she first chooses it in a standard chess game on Lichess. Entries in Panel B are difference-in-differences estimates of the memory premium. Columns 1B and 2B report results from a standard two-way fixed effects estimator with player-by-move and number-of-Chess960-games fixed effects, applied to the sample of all available moves in players' first fifty games of Chess960 as White. Columns 3B and 4B implement the estimator of Borusyak et al. (2024). The sample in these columns corresponds to all player-move combinations in which the move enters the player's memory either after her first game of Chess960 or not at all. Columns 5B and 6B implement the stacked difference-in-differences estimator of Cengiz et al. (2019). Each stack consists of moves that enter a given player's memory at a particular point in time as well as moves that remain out of the same player's memory over the entire sample period. In these specifications, all identifying variation comes from within-player comparisons of treated and never-treated moves over time. All estimates are scaled to correspond to the percentage-point change in choice probability associated with a one-unit increase in the respective regressor. Standard errors are clustered by player and shown in parentheses.

It is noteworthy, though, that our causal estimates of the memory premium are an order of magnitude smaller than those in Table 2. This is not surprising. On average, the moves that identify the estimates in Table 4 enter players' memories between their 16th and 17th game of Chess960 as White. By this time, the average memory premium has already decreased by about half. Moreover, being played for the first time in a standard chess game on Lichess constitutes a very weak treatment. For comparison, if we restrict attention to the 16th and 17th game of Chess960 and use the regression model in col. (6) of Table 2 to estimate the memory premium for moves that had been played only once before in regular chess, then we obtain a point estimate of 0.43 p.p. (with a standard error of 0.17). There is, therefore, no contradiction between the causal estimates in Table 4 and our estimates of the average memory premium in Table 2.

7. Mechanism

To establish that the memory premium is not specific to our particular empirical setting and to shed light on the mechanism behind it, we conducted a pre-registered laboratory experiment.[22]

7.1. *Experimental Design*

As explained in Section 2, there are at least two classes of mechanisms that may give rise to the memory premium. The first one involves utility maximization coupled with subjective differences in the expected utility of alternatives. According to this explanation, individuals assign higher expected utility to in- than to out-of-memory options—even when the latter are objectively just as good. Subjective differences in expected utility might arise because DMs form expectations about value by extrapolating from a familiar environment in which in-memory alternatives tend to perform well, or because DMs are risk-averse and hold more precise beliefs about alternatives that they have previously encountered.

The second class of mechanisms involves deviations from full maximization. According to this competing explanation, DMs do not necessarily consider all of the available alternatives. Instead, they rely on choice procedures that prioritize in- over out-of-memory options. One example of such a procedure is maximization from consideration sets that disproportionally include in-memory alternatives. Another example is satisficing, whereby the DM evaluates alternatives sequentially and selects the first one that is "good enough." Such a sequential choice procedure would give rise to a memory premium if the order of consideration favors in-memory alternatives.

[22]The preregistration for this experiment is available in the OSF Registries at `https://osf.io/8tkd4`. All analyses in the main text follow the preregistration. Additional analyses that were not pre-registered are presented in Appendix Tables AT.9–AT.13.

In order to distinguish between potential explanations for the memory premium, we now describe a new experimental design that isolates the second type of mechanism. In this design, we shut down the subjective utility channel. In addition, we make the order in which the available alternatives are being considered observable.

Setup In our experiment, subjects complete twenty rounds of a simple choice task. We describe the task as choosing one out of six "companies" for a hypothetical investment. Each company is identified by a unique name and color, and appears in only one round. Subjects' bonus payments are determined by the value of the company they choose in one randomly selected round. A company's value depends on its "profit" and "growth":

$$\text{Value} = \text{Profit} \times \text{Growth Factor}.$$

Profit levels are integers between 2 and 10, while growth factors are two-digit decimals between 0.5 and 1.5. Both are distributed uniformly and independently across companies, all of which is known to participants. Note, in this setup, a company's profit corresponds to its expected value to the participant prior to learning the growth factor.

Each round of the choice task consists of two stages, separated by a brief wait phase. In the first stage, which we call the "research stage," subjects are shown three of the six companies in their choice set, together with information on these companies' "number of stores," "revenue per store," and "total expenses." Subjects know that which companies appear in the research stage is entirely random. Based on the provided information, subjects are instructed to calculate the three companies' profits:

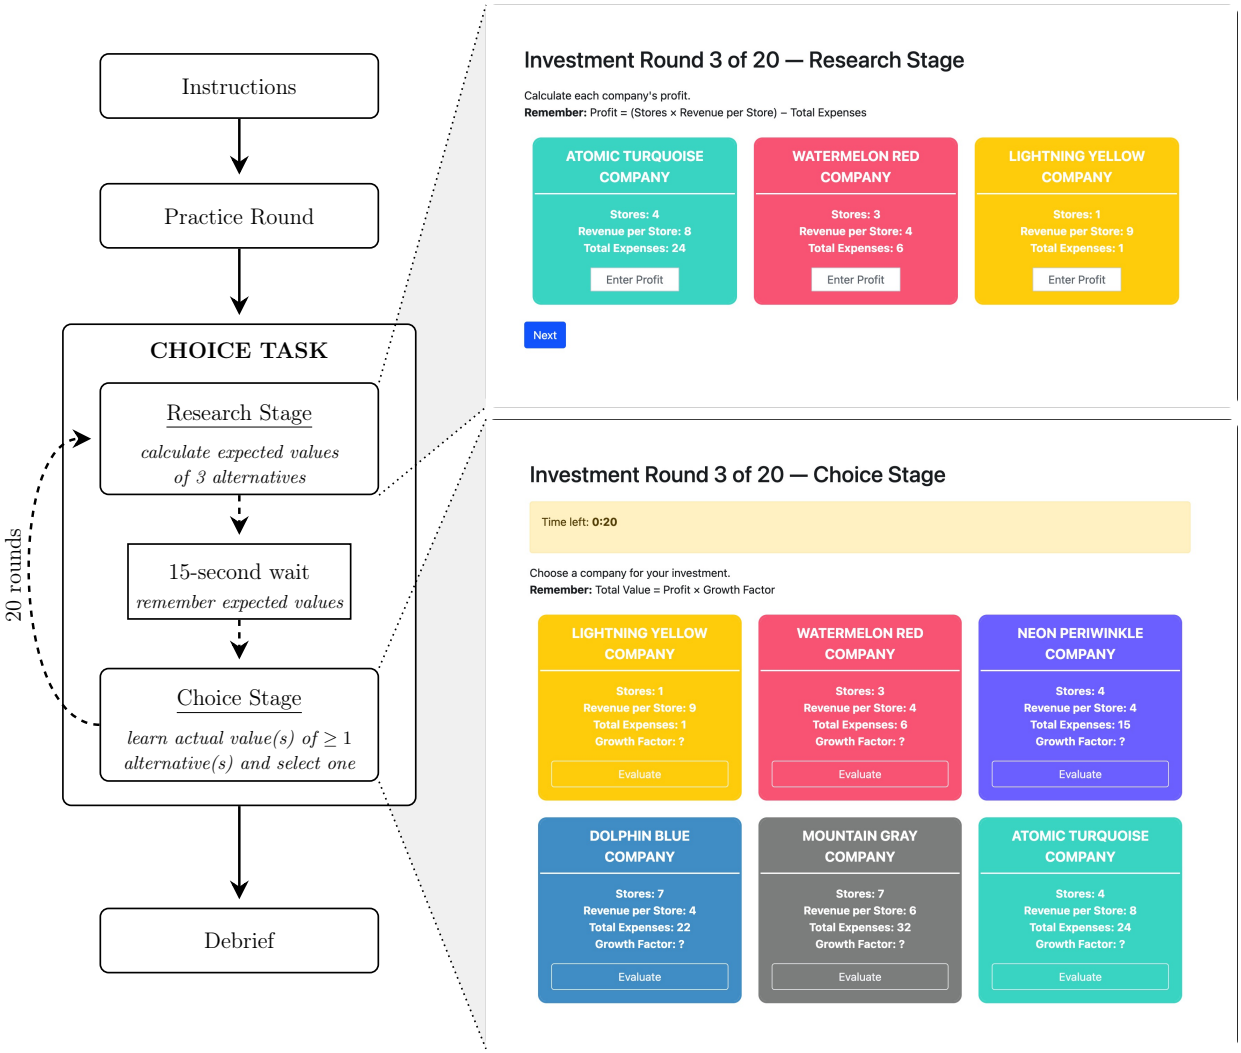
$$\text{Profit} = (\text{Stores} \times \text{Revenue per Store}) - \text{Total Expenses}.$$

Upon correctly calculating profits, subjects wait for fifteen seconds before proceeding to the "choice stage." In the choice stage, subjects are shown all six companies simultaneously.[23] At this point, subjects have enough information to calculate every company's profit, but they do not yet know the relevant growth factors. To see the growth factor of a particular company, subjects need to click on that company's "evaluate" button. Subjects may elicit the growth factors of as many companies as they wish, and they can do so in any order. The only constraints are that a company must be evaluated before it can be chosen, and that a choice must be made before the time limit for the choice stage expires. The time limit varies between ten and forty-five seconds, and is distributed uniformly across rounds.[24]

[23]The positioning of companies on the screen is random.

[24]Participants do not to make a timely choice in about 3.9% of problems. The analysis in the main text excludes all choice tasks for which this is the case. Appendix Table AT.10 shows that the results are robust

Figure 14: Experimental Design



Notes: Figure shows a high-level summary of our experimental design, together with screenshots from each stage of the choice task. The upper screenshot depicts the "research stage" of each round, in which participants are asked to calculate and memorize the expected value of three alternatives. The lower screenshot depicts the "choice stage." In the choice stage of a given round, subjects click on alternatives' "evaluate button" to reveal additional information about the respective option before making a final choice. The choice set always includes the three alternatives that subjects encountered during the "research stage" of the same round, as well as three new ones. In either stage, the position of alternatives on the screen is randomly determined.

Design Rationale Figure 14 provides a high-level summary of the experimental design, together with screenshots from both stages of the choice task. We settled on this design because we wish to isolate the role of the procedural mechanism in generating the memory premium. By *inducing* values and beliefs, our design eliminates any potential differences in the subjective utility of alternatives. The subjective utility channel, therefore, cannot account for the emergence of a memory premium in the experiment. In fact, given that an alternative

to including these observations. In addition, we have conducted a follow-up experiment in which we did not impose a binding time constraint. Again, the results are qualitatively equivalent (see Appendix Table AT.12).

must be evaluated before it can be chosen, participants face no uncertainty about the value of the selected option.

Another key feature of the design is that it allows us to monitor the choice process. Building on Johnson et al. (2002), we track which alternatives participants evaluated and in what order. We can, therefore, examine whether in-memory alternatives receive priority in consideration.

Finally, our two-stage design randomly induces memories of some alternatives in the choice set but not others. By instructing participants to remember the alternatives in the research stage and by restricting the delay before the choice stage to only fifteen seconds, we ensure that subjects are able to retrieve and rely on the relevant memories—if they wish. The results below should thus be interpreted as speaking to the memory premium under idealized conditions, when imperfect recall is unlikely to be a first-order issue.

Implementation The experiment took place on a custom-built website on July 12, 2025. We recruited 501 participants from Prolific, promising a $7.50 base payment plus a performance-based bonus of up to $15. The median participant earned a total of $16.18 and spent about 34.5 minutes on the experiment. For additional details and for a copy of the instructions to participants, see Appendix E.[25]

7.2. *Experimental Results*

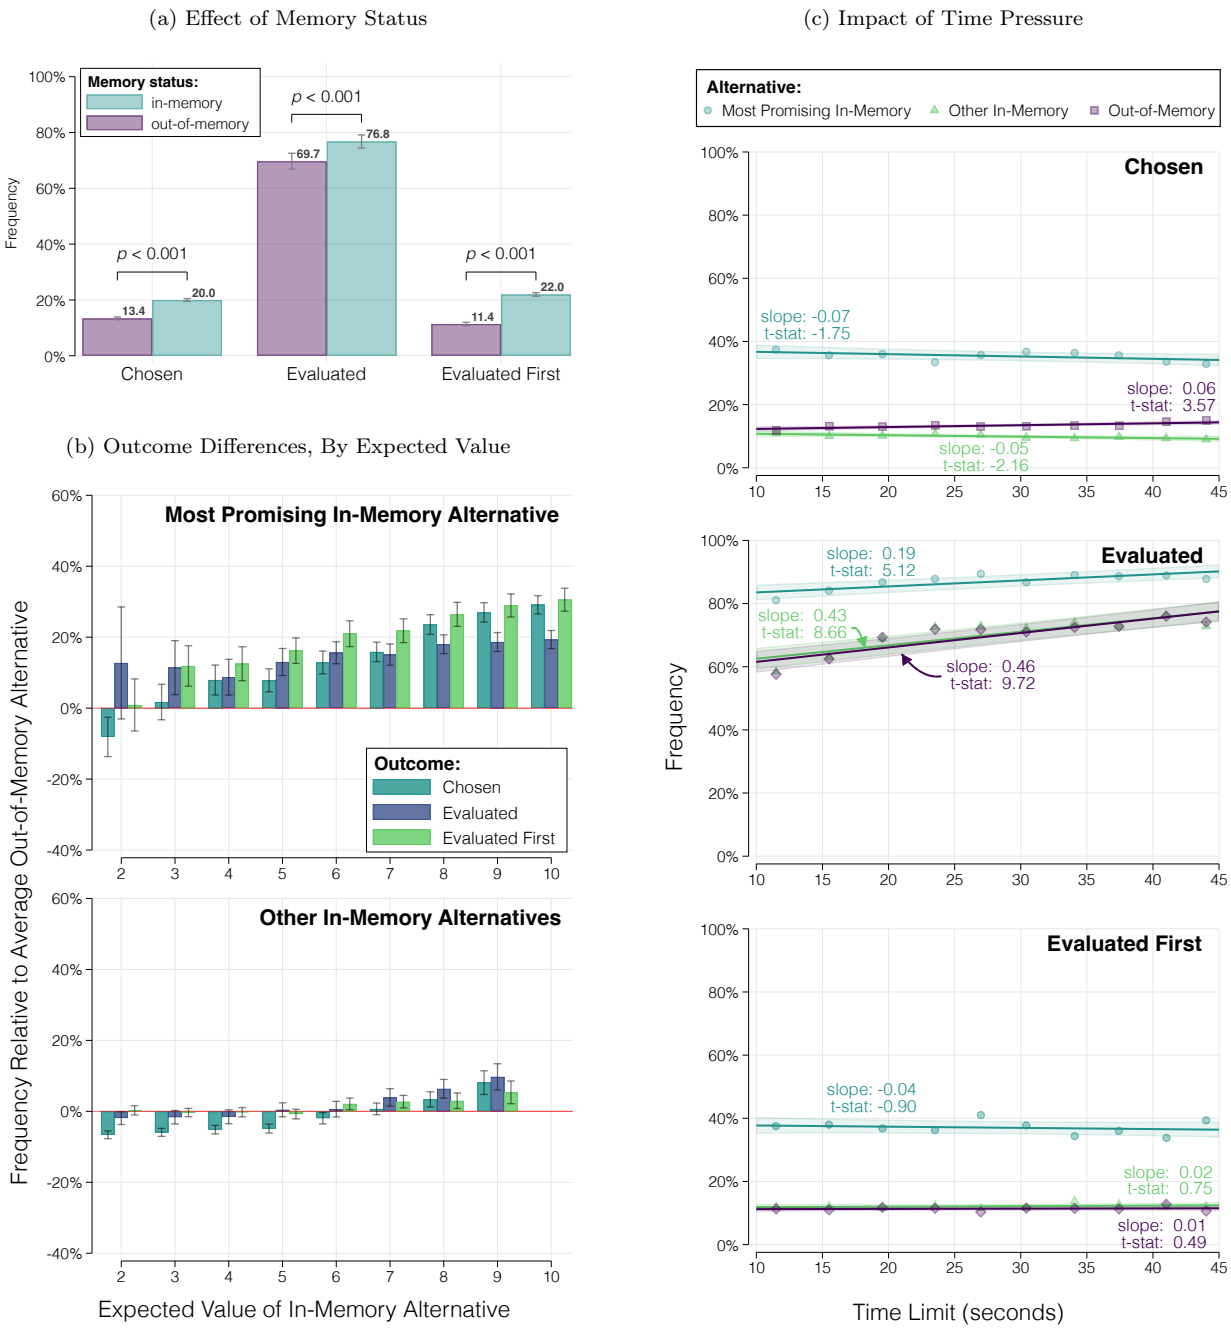
Figure 15 visualizes the results of the experiment. The top-left panel shows that in-memory alternatives are, on average, 6.6 p.p. more likely to be chosen than out-of-memory ones, even though there are no systematic differences in either expected or realized values. Relative to a mean choice frequency of 16.7%, the memory premium in this environment is economically large.

The top-left panel also shows that, on average, in-memory options are 7.1 p.p. more likely to be evaluated than out-of-memory ones; and they are 10.6 p.p. more likely to be evaluated first. In other words, subjects are not only more likely to consider alternatives that they have previously encountered, but they tend to start by considering these alternatives.

The results in the bottom-left panel of Figure 15 help to shed additional light on the choice procedure. We observe that the effects in the upper-left panel are driven by the "most promising" in-memory alternative, i.e., by the one with the highest expected value prior to learning the relevant growth factor. Relative to the *average* out-of-memory option, the most promising in-memory alternative is significantly more likely to be chosen, more likely to be evaluated, and more likely to be evaluated first. This is not the case for the two other

[25]More than 98% of participants report that the instructions they received were clear. For summary statistics on the participant population and selected outcomes, see Appendix Table AT.9

Figure 15: Experimental Results



Notes: Figure presents results from our experiment. Panel (a) compares mean outcomes for alternatives that appeared in the research stages of the choice task and were, therefore, in-memory with those that were not. Panel (b) compares mean outcomes of in-memory alternatives with a given expected value to the *average* out-of-memory alternative. The upper exhibit in this panel restricts attention to the in-memory alternative with the highest expected value. The lower exhibit focuses on the remaining two in-memory options. Panel (c) shows how outcomes depend on the time constraint in the choice stage. In doing so, it distinguishes between the in-memory alternative with the highest expected value, the remaining two in-memory alternatives, and out-of-memory options. Error bars and shaded regions correspond to 95%-confidence intervals, accounting for clustering by participant. For results that only compare in- and out-of-memory options with the same expected value and that further distinguish between the second- and third-most promising in-memory alternative, see Appendix Tables AT.10 and AT.11.

in-memory alternatives. To the extent that participants favor these alternatives over the average out-of-memory option, the estimated differences are much smaller.[26]

Importantly, even when the expected value of the most promising in-memory alternative is only three, four, or five, subjects tend to evaluate it first. They are also more likely to choose it than the average out-of-memory option. Given that the average value of the latter is six, subjects could, in expectation, do better by choosing a random out-of-memory option instead.[27]

The right panel of Figure 15 investigates how the observed behavior depends on the time constraint. Three patterns stand out: (i) Subjects' tendency to evaluate the most promising in-memory alternative first does not depend on the time limit. (ii) When given more time, subjects evaluate a greater number of in- and out-of-memory alternatives, but especially the latter. (iii) They thus become somewhat more likely to ultimately select an out-of-memory option. Importantly though, the choice frequency of the in-memory alternative with the highest expected value continues to exceed that of out-of-memory options by nearly twenty percentage points, irrespective of the time limit.[28]

Appendix Table AT.13 presents additional evidence from response times. The evidence therein indicates that participants complete the choice stage 8–13% faster when they select an in- rather than an out-of-memory alternative. Between 33% and 46% of this difference is explained by fewer evaluations.

7.3. *Discussion*

Based on these findings, we make three observations about the choice procedure. First, the data are inconsistent with any model of choice that treats in- and out-of-memory options symmetrically. Even if we allow for the possibility of mean-zero computational errors, such models imply that there should not be any difference in the choice frequencies of in- and

[26]The results in panel (b) of Figure 15 compare in-memory options with a particular expected value to the *average* out-of-memory alternative in order to condition on participants' information at the beginning of the choice stage. For results that compare in- and out-of-memory alternatives with the same expected value and that distinguish between the second- and third-most promising in-memory options, see Appendix Tables AT.10 and AT.11.

[27]Participants make small optimization errors, even conditional on which alternatives they consider. For instance, the average value of the chosen alternative equals about 89% of the maximum among all evaluated options. Optimization failures are positively correlated with the number of evaluations and negatively correlated with deliberation time (see Appendix Table AT.13), which suggests that these mistakes are partly due to imperfect recall and computational errors.

[28]In Appendix Table AT.11, we show that this finding continues to hold when we condition on both the time limit and alternatives' expected value. As an additional robustness check, we conducted a follow-up experiment without a binding time constraint in the choice stage ($N = 201$). The results from this experiment demonstrate that subjects prioritize the most promising in-memory alternative even when they do not have to make a decision quickly (see Appendix Table AT.12).

out-of-memory options. After all, there are no systematic differences in any other dimension, such as expected value. Special cases of symmetric models include full maximization, random search, and maximization from consideration sets.

For a model to explain the first-order patterns in the data, it needs to favor in-memory options. For instance, a sequential search procedure that prioritizes in-memory alternatives can explain why these alternatives are more likely to be evaluated and, ultimately, chosen. Maximization from consideration sets can explain our results if inclusion in the consideration set is a function of memory status. Although such modifications are theoretically trivial, they are empirically relevant.

Our second observation is that the most promising in-memory alternative occupies a special place in the decision-making process *even when its expected value is smaller than that of a randomly selected out-of-memory option*. This finding is difficult to reconcile with rational search. To understand why, consider situations in which participants chose the most promising in-memory alternative despite its expected value being five or less. Although the chosen alternatives do end up being positively selected, their average realized value in these instances is only 5.20, with a 95%-confidence interval of $[5.02, 5.39]$. By picking a random out-of-memory alternative, however, subjects could have guaranteed themselves an expected payoff of six. Our results, therefore, hint at a purely procedural component of the decision-making process. At the same time, we observe that the most promising in-memory alternative is prioritized to a greater extent when it is objectively better, suggesting that the observed behavior may be boundedly rational.

Our third observation is that the second- and third-most promising in-memory alternatives are also favored in the decision-making process, although not as much as the in-memory option with the highest expected value. To see this, note that the average expected value of these two alternatives is less than five, while that of out-of-memory ones equals six. Yet, as can be seen in panel (c) of Figure 15, both kinds of alternatives are nearly equally likely to be evaluated first.

Taken together, our results point to a procedural explanation for the memory premium. In-memory alternatives are not only more likely to be chosen than out-of-memory ones, but they are more likely to be considered for choice, and more likely to be considered first. More broadly, our findings suggest that individuals rely on their memory to prioritize options that warrant closer examination, albeit in a less-than-perfectly rational way.

To be clear, our experimental results do not imply that the subjective utility channel is quantitatively unimportant in other settings. To the extent that in-memory alternatives are assigned higher subjective expected utility than out-of-memory ones, such differences would further reinforce the memory premium.

8. Concluding Remarks

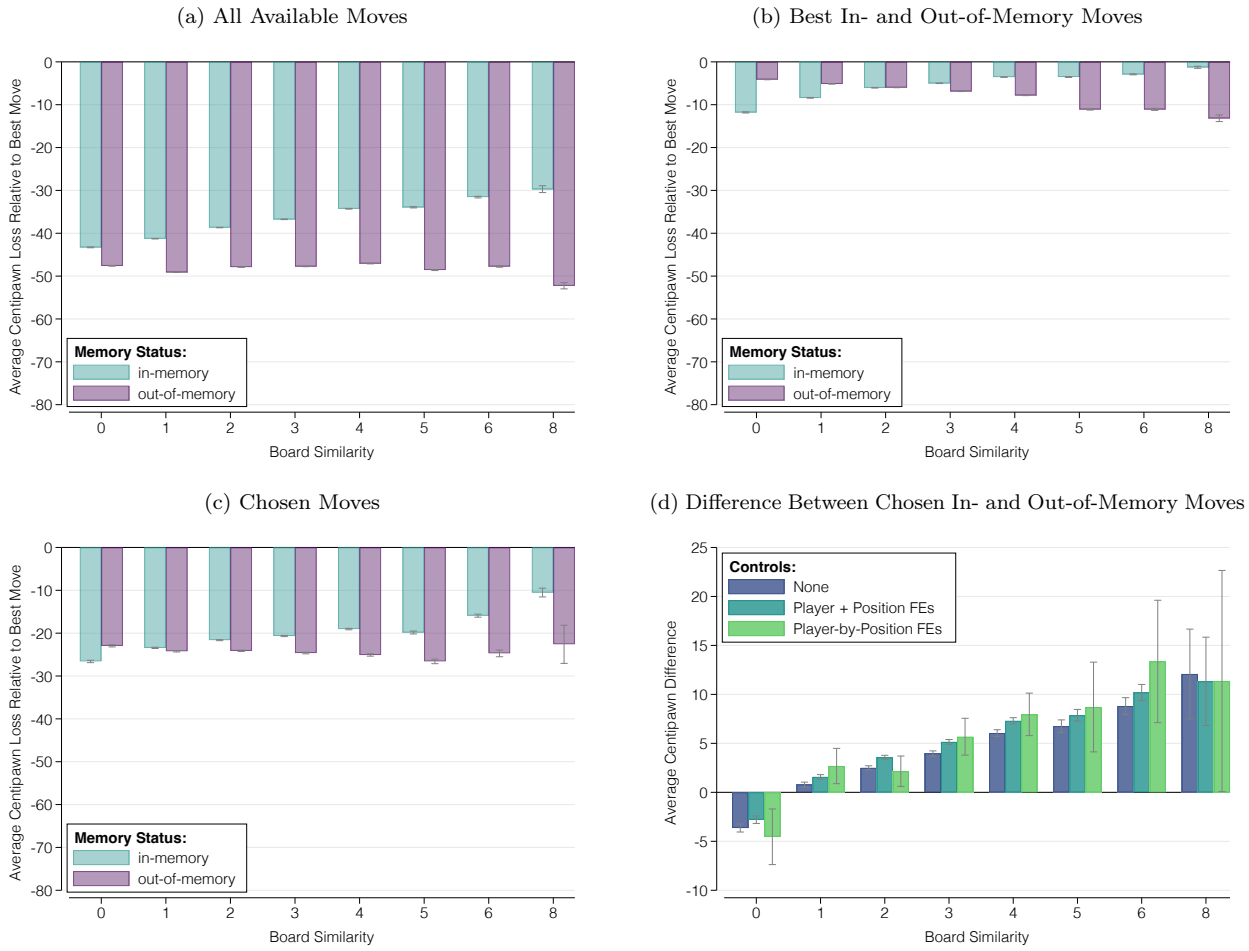
Memory is a key ingredient in human decision-making, allowing us to use past experiences to guide future behavior. This paper documents that, in unfamiliar environments, decision-makers are more likely to choose alternatives that are retrievable from memory than those that are not, even when the latter are objectively better. We dub this phenomenon the "memory premium," and provide evidence that it is associative, subject to interference and repetition effects, and that it decays over time. While some of these comparative statics are also consistent with alternative explanations that do not directly appeal to memory, it is not obvious that there exists another explanation for the entirety of our observational and experimental results. Taken together, our findings imply that the ease with which past experiences come to mind plays an important role in shaping choice behavior, at least in part because decision makers are more likely to seriously consider in-memory options.

Taking a step back, we can ask whether it makes sense to rely on memory as a guide for decision-making in unfamiliar environments. A key benefit of memory-based heuristics is that they economize on cognitive costs. Suppose that there are N unfamiliar environments and that the cost of analyzing every new environment to find an environment-specific optimal strategy is c. The cost of computing an optimal mapping from environments to strategies is thus cN. An extreme way to economize on cognitive costs would be to use a constant mapping instead, i.e., a strategy that does not change with the environment. Relying entirely on memory rather than careful analysis is an example of such a mapping. Of course, there is a wide range of intermediate cases. For example, a DM may utilize memory when the decision problem is sufficiently similar to the ones that she encountered in the past but engage in deeper analysis when it is not. In this sense, drawing on memory to make decisions in unfamiliar environments can be viewed as a cognitive shortcut.

But does such a shortcut yield good outcomes? While a general answer is beyond the scope of this paper, we note that the performance of any memory-based heuristic will, at least in part, depend on the degree of similarity between familiar and unfamiliar environments. Focusing on the case of Chess960, the top two panels of Figure 16 examine the average quality of opening moves, by memory status. The top-left panel shows the average centipawn rating of in- and out-of-memory moves relative to the best available alternative in a given starting position. In-memory alternatives are, on average, better than out-of-memory ones, with the gap widening as the starting position becomes more similar to the standard board. As a result, a naïve decision maker could, in expectation, improve upon random choice by randomly choosing an in-memory alternative. The magnitude of the improvement would be especially large in starting positions that closely resemble the standard one.

The top-right panel of Figure 16 compares the *best* available in- and out-of-memory moves

Figure 16: Quality of Moves, by Memory Status



Notes: Panel (a) compares the average Stockfish rating of all available in- and out-of-memory moves relative to the best move in a particular starting position. Negative values thus correspond to losses relative to the optimum. Panel (b) shows the average Stockfish rating of players' highest-rated in- and out-of-memory moves in a particular starting position relative to the best opening move in that position. Panel (c) shows the average Stockfish rating of the moves that players actually choose, by memory status and relative to the best available alternative. Panel (d) presents estimates of the difference in Stockfish ratings between chosen in- and out-of-memory moves, controlling for different sets of fixed effects. All panels differentiate according to the similarity of the starting position in Chess960 with that in standard chess. The sample always consists of moves in players' first fifty Chess960 games as White. Error bars correspond to 95%-confidence intervals, accounting for clustering by player.

in a given starting position. For high and intermediate levels of board similarity, the players in our data would, on average, sacrifice very little if they maximized only over the consideration set of in-memory alternatives. Only in the most dissimilar board positions would such a choice procedure result in a large expected loss relative to the first best.

The bottom two panels focus on the opening moves that were actually chosen. The panel on the bottom-left presents raw data, while the one on the right shows estimates of the difference in quality between in- and out-of-memory moves after controlling for player and starting position fixed effects. Two findings emerge. First, conditional on memory status, chosen moves are far better than average but significantly worse than the best available

alternative, suggesting that choices in this setting are deliberate yet imperfect. Second, unless the starting position is highly dissimilar from the standard one, players do better when they choose in- rather than out-of-memory moves. This pattern is not only borne out in the raw data, but it also holds conditional on player-by-starting-position fixed effects, i.e., when we compare choice situations in which the same player encountered the same starting position more than once. Only in the most dissimilar starting positions do we observe that players do better when they choose moves that are out-of-memory.

To help interpret the previous observation, note that although in-memory moves are, on average, better than out-of-memory ones for all levels of board similarity (see panel (a)), in dissimilar starting positions the best out-of-memory alternatives tend to be better than the best in-memory alternatives (see panel (b)). As a consequence, players can realize nontrivial improvements by carefully considering *all* available options.

In sum, the evidence above is consistent with the idea that relying on memory is boundedly rational. It can reduce cognitive costs while yielding favorable outcomes in environments that are not too dissimilar.

References

Afrouzi, H., S. Kwon, A. Landier, Y. Ma, and D. Thesmar (2023). Overreaction in Expectations: Evidence and Theory. *Quarterly Journal of Economics 138*(3), 1713–1764.

Azeredo da Silveira, R., Y. Sung, and M. Woodford (2024). Optimally Imprecise Memory and Biased Forecasts. *American Economic Review 114*(10), 3075–3118.

Azeredo da Silveira, R. and M. Woodford (2019). Noisy Memory and Over-Reaction to News. *AEA Papers and Proceedings 109*, 557–561.

Baddeley, A., M. Eysenck, and M. Anderson (2009). *Memory* (3rd ed.). Hove, UK: Psychology Press.

Bordalo, P., G. Burro, K. Coffman, N. Gennaioli, and A. Shleifer (2024). Imagining the Future: Memory, Simulation, and Beliefs. *Review of Economic Studies*, forthcoming.

Bordalo, P., J. Conlon, N. Gennaioli, S. Kwon, and A. Shleifer (2022). Memory and Probability. *Quarterly Journal of Economics 138*(1), 265–311.

Bordalo, P., N. Gennaioli, and A. Shleifer (2020). Memory, Attention and Choice. *Quarterly Journal of Economics* (3), 1399–1442.

Borusyak, K., X. Jaravel, and J. Spiess (2024). Revisiting Event-Study Designs: Robust and Efficient Estimation. *Review of Economic Studies 91*(6), 3253–3285.

Calonico, S., M. Cattaneo, and R. Titiunik (2014). Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs. *Econometrica 82*(6), 2295–2326.

Cengiz, D., A. Dube, A. Lindner, and B. Zipperer (2019). The Effect of Minimum Wages on Low-Wage Jobs. *Quarterly Journal of Economics 134*(3), 1405–1454.

Charles, C. (2022). Memory and Trading. mimeographed, London School of Economics.

Charles, C. (2024). Memory Moves Markets. mimeographed, London School of Economics.

Chen, D., M. Schonger, and C. Wickens (2016). oTree - An Open-Source Platform for Laboratory, Online and Field Experiments. *Journal of Behavioral and Experimental Finance 9*, 88–97.

D'Acunto, F., U. Malmendier, J. Ospina, and M. Weber (2021). Exposure to Grocery Prices and Inflation Expectations. *Journal of Political Economy 129*(5), 1615–1639.

Enke, B., F. Schwerter, and F. Zimmermann (2024). Associative Memory, Beliefs and Market Interactions. *Journal of Financial Economics 157*, 103853.

Gilboa, I. and D. Schmeidler (1995). Case-Based Decision Theory. *Quarterly Journal of Economics 110*(3), 605–639.

Gilboa, I. and D. Schmeidler (2001). *A Theory of Case-Based Decisions*. Cambridge, UK: Cambridge University Press.

Graeber, T., C. Roth, and F. Zimmermann (2024). Stories, Statistics, and Memory. *Quarterly Journal of Economics 139*(4), 2181–2225.

Hartzmark, S., S. Hirshman, and A. Imas (2021). Ownership, Learning, and Beliefs. *Quarterly Journal of Economics 136*(3), 1665–1717.

Jiang, Z., H. Liu, C. Peng, and H. Yan (2024). Investor Memory and Biased Beliefs: Evidence from the Field. NBER Working Paper No. 33226.

Johnson, E., C. Camerer, S. Sen, and T. Rymon (2002). Detecting Failures of Backward Induction: Monitoring Information Search in Sequential Bargaining. *Journal of Economic Theory 104*(1), 16–47.

Kahana, M. (2012). *Foundations of Human Memory*. New York, NY: Oxford University Press.

Malmendier, U. and S. Nagel (2011). Depression Babies: Do Macroeconomic Experiences affect Risk Taking? *Quarterly Journal of Economics 126*(1), 373–416.

Malmendier, U. and S. Nagel (2016). Learning from Inflation Experiences. *Quarterly Journal of Economics 131*(1), 53–87.

Malmendier, U. and J. Wachter (2024). Memory of Past Experiences and Economic Decisions. In M. Kahana and A. Wagner (Eds.), *Oxford Handbook of Human Memory, Volume II: Applications*. New York, NY: Oxford University Press.

Manzini, P. and M. Mariotti (2014). Stochastic Choice and Consideration Sets. *Econometrica 82*(3), 1153–1176.

Melton, A. (1963). Implications of Short-Term Memory for a General Theory of Memory. *Journal of Verbal Learning and Verbal Behavior 2*(1), 1–21.

Mullainathan, S. (2002). A Memory-Based Model of Bounded Rationality. *Quarterly Journal of Economics 117*(3), 735–774.

Nagel, S. and Z. Xu (2022). Asset Pricing with Fading Memory. *Review of Financial Studies 35*(5), 2190–2245.

Salant, Y. and J. Spenkuch (2024). Complexity and Satisficing: Theory with Evidence from Chess.

Review of Economic Studies, forthcoming.

Scharnagl, R. (2004). *Fischer-Random-Schach (FRC / Chess960): Die revolutionäre Zukunft des Schachspiels*. Hamburg: Books on Demand GmbH.

Simon, H. (1955). A Behavioral Model of Rational Choice. *Quarterly Journal of Economics 69*(1), 99–118.

Tversky, A. and D. Kahneman (1973). Availability: A Heuristic for Judging Frequency and Probability. *Cognitive Psychology 5*(2), 207–232.

Wachter, J. and M. Kahana (2024). A Retrieved-Context Theory of Financial Decisions. *Quarterly Journal of Economics 139*(2), 1095–1147.

Woodford, M. (2020). Modeling Imprecision in Perception, Valuation, and Choice. *Annual Review of Economics* (12), 579–601.

Online Appendix

Contents

A Proof of Observation 1 **3**

B Data Appendix **3**

C Additional Results and Robustness Checks **3**

D Importance of Opening Strategies **5**

E Experimental Details **6**

Experimental Screens **8**

Appendix Figures **18**

Appendix Tables **22**

List of Figures

AF.1 Distribution of Starting Positions in Chess960 18
AF.2 Choice Frequency of Opening Moves in Standard Chess 18
AF.3 Estimated Memory Premia, Controlling for the Set of Other In-Memory Moves 19
AF.4 Decay in the Memory Premium, Based on Donut Definitions of Memory 19
AF.5 Frequency Effects, By Prior Success 20
AF.6 Dynamics of the Memory Premium, Based on Prior Moves in Chess960 20
AF.7 Density of Observations around RD Threshold 21

List of Tables

AT.1 Memory Premium, Controlling for All Other In-Memory Moves 22
AT.2 Memory Premium, by Quality of Move in a Particular Starting Position 23
AT.3 Memory Premium, Flipped Standard Board 24
AT.4 Memory Premium, Boards with Corner Royalty 25
AT.5 Memory Premium, 1-Hour Donut-Definition of Memory 26
AT.6 Memory Premium, 1-Day Donut-Definition of Memory 26
AT.7 Memory Premium, 1-Week Donut-Definition of Memory 27

AT.8 Memory Premium, 1-Month Donut-Definition of Memory 27
AT.9 Summary Statistics for Experimental Data 28
AT.10 Experimental Results, Including Choice Failures 29
AT.11 Experimental Results, by Ranking of Alternatives in the Research Stage 30
AT.12 Results from Robustness Experiment 31
AT.13 Ancillary Experimental Results 32

Appendix A: Proof of Observation 1

By continuity, it suffices to show that $C(S, a \mid M) > C(S, b \mid M)$ when $V(b, E) = V(a, E)$. Given Definition 1, we have that

$$C(S, a \mid M) - C\left(S, a \mid M^{-a}\right) = \Delta(S, a \mid M, a) > \Delta(S, b \mid M, a) = C(S, b \mid M) - C\left(S, b \mid M^{-a}\right).$$

The observation follows because monotonicity and $V(b, E) = V(a, E)$ imply that $C(S, a \mid M^{-a}) = C(S, b \mid M^{-a})$.

Appendix B: Data Appendix

The chess data in this paper were furnished by `lichess.org`. Every day, Lichess hosts several million games of standard chess and thousands of games of Chess960. In the summer of 2021, we contacted the founder of Lichess and purchased a database extract covering *all* standard and variant games that were played on the platform between January 2013 and June 2021.[1] The raw data were provided to us in the human-readable PGN format and include basic facts about each game (including players' usernames and ratings, date and time of the game, time controls, ultimate outcome, etc.), as well as the exact sequence of moves. The total size of the raw data is about 6.9TB. In order to reconstruct users' history of play, we process these data using automated scripts and extract all opening moves of every registered player.

The sample for our main analysis restricts attention to the opening moves of 147,357 players in their first game of Chess960 on the platform. In order to be included in this sample, users must have played at least 20 standard chess games prior to their first game of Chess960. This restriction helps to ensure a minimal level of accuracy in approximating players' memory. We further require that their first game of Chess960 was rated, and that they played White. We focus on moves as White because we want to analyze their initial choices in an unfamiliar decision problem, and we require that the game be rated in order to ensure that these are real-stakes decisions.

In the second part of our analysis, we study the opening moves of all registered users who played a total of at least fifty games of Chess960 as White, regardless of which color they were assigned in their first game and whether that game was rated. There are 16,731 such players in our data.

In order to measure the quality of opening moves, we use the Stockfish chess engine to rate every feasible move in every possible starting position. All ratings are calculated using a single-threaded instantiation of Stockfish version 16.1, with depth set to 40, 4,096MB hash, the default neural net, and access to the 3-,4-,5-, and 6-man Syzygy tablebases.

Appendix C: Additional Results and Robustness Checks

This appendix reports additional results and robustness checks. The order in which these are discussed corresponds to the order in which they are referenced in the main text.

[1] Lichess regularly releases database extracts for rated chess games at `https://database.lichess.org`. These publicly available extracts do not, however, include casual games.

Appendix Figure AF.1 plots the frequency of every possible starting position in Chess960 in our main sample of players' first game of Chess960 as White. The reported χ^2-test indicates that the observed frequencies are consistent with the null hypothesis of uniform random assignment.

Appendix Figure AF.2 shows the frequency with which the players in the main sample chose each opening move in standard chess games on Lichess prior to their first game of Chess960.

Appendix Table AT.1 replicates Table 2 in the text controlling for *which* other moves players hold in memory. Specifically, the regression models in this table replace the move-by-starting-position-by-size-of-memory fixed effect in Table 2 with a fixed effect for the combination of the move, starting position, and the set of other in-memory moves.

Appendix Figure AF.3 replicates Figure 7 in the text controlling for *which* other moves players hold in memory. The estimates in this figure are based on a regression model that replaces the move-by-starting-position-by-size-of-memory fixed effect in the specification in col. (6) of Table 2 with a fixed effect for the combination of the move, starting position, and the set of other in-memory moves.

Appendix Table AT.2 reports estimates of the memory premium, distinguishing between all starting positions, as well as positions in which the move is and is not close to optimal (i.e., within 20 centipawns of the optimal move in a given position). All estimates in this table are based on the regression model and sample in col. (6) of Table 2, allowing for heterogeneity in the memory premium across opening moves.

Appendix Tables AT.3 and AT.4 provide estimates of the memory in the "flipped" standard board position (i.e., when only the king and queen sit on opposite squares) and in board positions in which the king and queen are placed near the corners of the board (i.e., on squares a1, b2, g1, or h1). The evidence in these tables helps to rule out that our findings are driven by players attempting to execute particular opening strategies.

Appendix Figure AF.4 and Appendix Tables AT.5–AT.8 respectively replicate the analyses in Figure 9(b) and Table 2 in the text based on "donut definitions" of memory. According to these definitions, a move is said to be in-memory if, and only if, it had been played at least once sufficiently long prior to a player's first game of Chess960.

Appendix Figure AF.5 replicates the analysis in Figure 9(a) in the text, differentiating between moves that players used more or less successfully in prior standard games on Lichess, as measured by the outcomes of the respective games. Specifically, we distinguish between moves with above- and below-median *player-specific* average payoffs in prior use in standard chess.

Appendix Figure AF.7 shows the density of observations around the threshold in the regression discontinuity design in Figure 13 and Table 4 in the text.

Appendix Table AT.9 presents descriptive statistics for our experimental data, including participant demographics provided by Prolific.

Appendix Tables AT.10 and AT.11 show that our experimental results are robust to controlling for alternatives' expected value, their display position, the time limit, as well as to including rounds in which participants did not submit timely choices. Per the pre-registration, our analysis in the main text excludes all observations from such rounds.

Appendix Table AT.12 reports results from the untimed robustness experiment described in Appendix E. Though smaller, the coefficients in this table are qualitatively similar to their counterparts in Figure 15 in the main text.

Appendix Table AT.13 reports correlations between deliberation times, the number of evaluated alternatives, and optimization errors in our main experiment.

Appendix D: Importance of Opening Strategies

The results in Section 4 demonstrate that first-time players of Chess960 are significantly more likely to choose opening moves that are in memory on account of having previously been used in regular chess. One possible—perhaps even natural—explanation for this behavior is that players simply execute familiar opening strategies. According to this explanation, players may choose ♙e4 not because it is in-memory but because they seek to establish control over the center of the board. Likewise, they may be more likely to play ♘f3 not because this move is in-memory but because the Zukertort Opening allows them to develop the king-side knight while maintaining strategic flexibility for subsequent moves. If individuals stick with opening strategies that are familiar to *them*, then this might create the *appearance* of a memory premium, even after controlling for move-by-starting-position fixed effects.

In order to rule out this potential explanation, we examine players' choices in starting position 534.[2] In this position, only the king and queen are swapped, making it a mirror image of the starting position in regular chess. If players are simply executing particular opening strategies, then, rather than choosing in-memory alternatives, they should now select the mirror counterparts of the respective moves. For example, a player who usually starts with the Zukertort Opening should now choose ♘c3 instead of ♘f3.[3] The results in Table AT.3, however, imply that this is not the case.

The coefficients in odd-numbered columns of this table show that, depending on controls, the average memory premium in starting position 534 equals about 3.1–3.8 p.p. Even-numbered columns run horse races between the memory status of a given move and that of its mirror counterpart (e.g.,

[2]For the official numbering of starting positions in Chess960, see Scharnagl (2004).

[3]The correlation between the Stockfish ratings of moves in the standard starting position and the ratings of their mirror counterparts in position 534 is 0.93. It is not exactly equal to one because Stockfish's internal heuristics do not directly exploit symmetry; and because castling paths in starting position 534 differ by one tempo. By contrast, the correlation between the same moves' ratings in the standard opening position and position 534 is only 0.33.

♘c3 ⇌ ♘f3). While we do observe that players are somewhat more likely to choose a particular move when its mirror is in-memory on account of having been played in regular chess, we observe that the coefficients on the memory status of the mirror move are statistically insignificant and much smaller than the respective estimates of the memory premium. We also observe that the estimated memory premium changes by less than 0.2 p.p. (6.4%) when we account for the memory status of the mirror move.

Moreover, Appendix Table AT.4 shows that the memory premium is quantitatively robust to restricting attention to starting positions in which the king and queen are placed near the corners of the board (i.e., on squares a1, b1, g1, or h1). Even in these positions, players tend to choose in-memory alternatives—especially center-pawn moves (i.e., ♙d3, ♙d4, ♙e3, and ♙e4). Based on these observations, we conclude that the memory premium is unlikely due to players simply executing familiar opening strategies. We also note that reliance on particular opening strategies cannot explain our experimental results in Section 7 of the main text.

Appendix E: Experimental Details

As explained in the text, our pre-registered main experiment ran on July 12, 2025. We recruited 501 participants via Prolific, promising a $7.50 base payment and a performance-based reward of up to $15. The experiment itself took place on a custom website that we built using oTree (Chen et al. 2016). To accommodate up to 50 simultaneous participants, we hosted the site on a "Performance-L" instance on Heroku.

The experiment had five phases: 1. Consent; 2. Instructions; 3. Practice Round; 4. Choice Task (2 stages, 20 rounds); 5. Debrief (2 questions).

The heart of the experiment consisted of twenty rounds of a simple choice task. As illustrated in Figure 14 in the text, the choice task had two stages separated by a 15-second wait. We described the task as choosing one out of six fictitious companies for a hypothetical investment. Each company is identified by a unique color and a corresponding name. No company appears in more than one round. Colors and names were generated by artificial intelligence software that was instructed to create a palette of 126 visually distinct colors with "snappy, memorable" names.

In the first stage of the choice task, which we refer to as the research stage, participants are shown three out of the six companies in their choice set, along with information that allows them to determine the profits of each of these alternatives through simple calculations. Profits correspond to the expected value of the alternative prior to the respective company's growth factor being revealed. Whether or not a particular alternative appears in the research stage is random. Participants further know that profits are round numbers between two and ten, and that any of these values is equally likely.

To proceed to the second stage of the choice task, which we call the choice stage, participants must correctly calculate the profits of all three companies, after which they are made to wait for fifteen seconds. In the choice stage, participants are shown all six companies in their choice set, along with

sufficient information to calculate the expected value of any alternative. The ordering of alternatives on the screen is random.

To determine an alternative's actual value, participants must elicit a final piece of information. To do this, they need to click on the respective company's Evaluate button. Clicking this button reveals the company's growth factor. It also enables another button that allows participants to select the respective alternative. Subjects may evaluate as many companies as they wish, in any order. The only constraint we impose is that a company must be evaluated before it can be chosen, and that a selection must be made before the time limit for that round expires.

Subjects had between ten and forty-five seconds to complete the choice stage. The time limit was randomized and distributed uniformly and i.i.d. across rounds. As indicated in the preregistration, the results in the main text do not consider rounds in which participants failed to make a timely choice, which occurred in about 3.9% of tasks. We obtain results that are qualitatively equivalent when we include these observations, coding none of the available alternatives as having been chosen (see Appendix Table AT.10).

As a robustness check, we conducted a follow-up experiment with 201 participants, in which we did not impose a binding time constraint.[4] The follow-up experiment ran on July 26, 2025, and involved no other changes. It was, therefore, not pre-registered. Results from the follow-up experiment are qualitatively similar and are presented in Appendix Table AT.12.

Subjects earned a performance-based bonus equal to the value of the alternative they selected in one randomly determined round. The median participant earned a total of $16.18 (incl. a $7.50 base payment) and spent about 34.5 minutes on the experiment.[5]

Appendix Table AT.9 presents descriptive statistics for our experimental data, including participant demographics provided by Prolific. On the following pages, we reproduce the text that was shown to participants, with horizontal lines demarcating screens.

[4]For technical reasons, the choice stage in the follow-up experiment ended after 10 minutes, which corresponds to more than twenty-five times the maximal time limit in the main experiment. In our view, subjects who fail to submit a choice within ten minutes are likely not paying attention to the task.

[5]These numbers refer to participants in the main experiment. Subjects in the follow-up experiments faced the same incentives, took about 45 minutes to complete the experiment, and earned $16.57, on average.

Experimental Screens

Research Study

Research Study: Prolific Choice Experiments (STU00224136)
Principal Investigators: Dr. Yuval Salant; Dr. Jörg Spenkuch
Supported By: This research is funded by Northwestern University.

Welcome! The purpose of this study is to better understand how memory influences decision-making. We are very grateful for your help!

To participate in this study you must be at least 18 years old and reside in the U.S. We also ask that you only participate if you are using a computer. Tablets and phones are not supported. Below we provide additional information about this study in order to help you decide whether you'd like to participate.

To get started, you need to provide your consent by pressing the PROCEED button at the bottom of this page.

What should I expect?
Your participation is voluntary. If you choose to participate, you will be asked to perform basic mathematical calculations and to choose between different hypothetical investment options. We estimate that it will take about 45 minutes to complete the study.

Will I be paid?
We will reward completion of this study with a base payment of $7.50, and a performance-based reward up to $15. Your performance-based reward will depend on your ability to identify good investment options.

Are there any risks?
We foresee little risk from participating in this study, and we do not guarantee that you will receive any benefits beyond the payments described above.

How will my information be used?
The information collected in this study will be exclusively used for research purposes. All data will be handled and stored in accordance with Northwestern University policy. There is minimal risk that participants might be identified from the information provided. The research team will take extensive precautions to keep all data secure in order to protect confidentiality. As part of this effort, your actual identity will remain unknown to the researchers conducting this study. The results of this research may be published, but only in anonymized form.

Who can I talk to?
If you have questions, concerns, or complaints, you can contact the Principal Investigators at

prolific@u.northwestern.edu. This research has been reviewed and approved by an Institutional Review Board (IRB) — an IRB is a committee that protects the rights of people who participate in research studies. You may contact the IRB by phone at +1 (312) 503-9338 or by email at irb@northwestern.edu if:

- Your questions, concerns, or complaints are not being answered by the research team.
- You cannot reach the research team.
- You want to talk to someone besides the research team.
- You have questions about your rights as a research participant.
- You want to get information or provide input about this research.

By proceeding to the next screen, you are consenting to participate in this study.

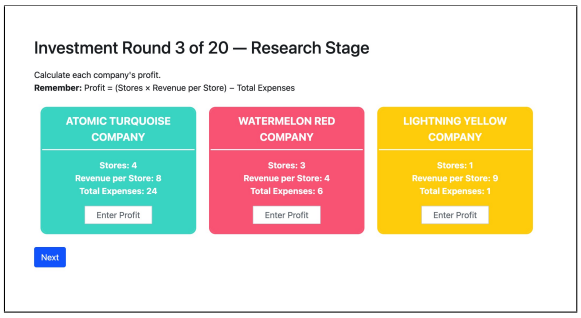
Instructions — Page 1 of 5

In this study, you will be asked to make hypothetical investment decisions. Each decision involves choosing one out of six fictitious companies. Prior to making a choice, you will be given the opportunity to memorize information that partially determines the value of three of the six companies.

There are twenty rounds in this study. In each round, you will make one investment decision. At the end of the study, we will pick one decision at random, and you will be paid the value of the company that you chose in that round. Your performance-based bonus, therefore, depends on the quality of your decisions.

To see an example of what exactly you're being asked to do in each investment round, please proceed to the next screen.

Instructions — Page 2 of 5



Each round in this study consists of three stages: a research stage, a wait phase, and a choice stage.

In the **research stage**, you will be given information on three companies. As in the screenshot above, the information on each company includes its number of stores, the revenue per store, and its total expenses.

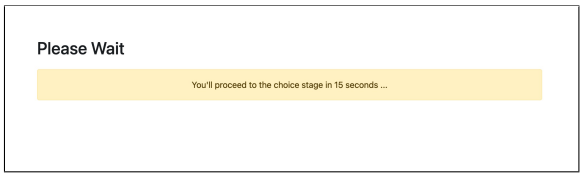
Based on this information, you'll be asked to calculate each company's profit. Profit is determined according to the following formula:

$$\text{Profit} = (\text{Stores} \times \text{Revenue per Store}) - \text{Total Expenses}$$

For example, applying this formula to the Lightning Yellow Company in the screenshot above, its profit equals 8 [= (1 × 9) - 1]. The profit of the Watermelon Red Company, however, is 6 [= (3 × 4) - 6].

You will not be able to proceed to the next screen until you have entered the correct profit numbers for all three companies. Profits will play an important role in the choice stage.

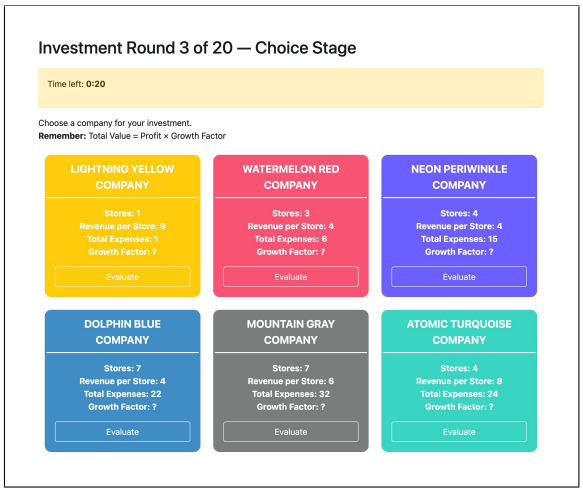
Instructions — Page 3 of 5



Every research stage is followed by a brief **wait phase**. The length of the wait is 15 seconds.

While you wait, you are being asked to remember the profit of each of the three companies that you encountered during the research stage.

Instructions — Page 4 of 5



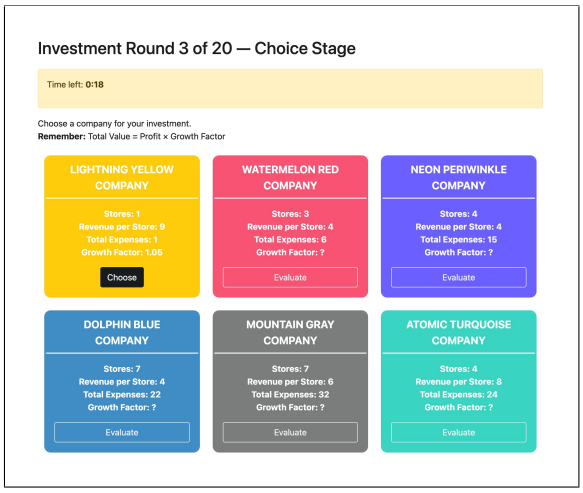
In the **choice stage** of each round, you are being asked to invest in one out of six companies, three of which you have already encountered during the research stage.

The total value of a given company depends on its profitability and growth. Companies that are growing rapidly have a *growth factor* greater than one, whereas companies that are likely to go out of business have a growth factor smaller than one. More specifically, the value of a given company is determined according to the formula:

$$\text{Value} = \text{Profit} \times \text{Growth Factor}$$

While you may remember the profit numbers for the companies that you encountered in the research stage, you will need to click on the Evaluate button in order to discover a company's growth factor.

Instructions — Page 5 of 5



In the example above, the growth factor of the Lightning Yellow Company is 1.05. Since the company's profit is 8 [= 1 × 9 - 1], its total value equals 8.4 [= 8 × 1.05].

You may click on the evaluate buttons of as many companies as you like. After evaluating a particular company, you can select that company for your investment by clicking on the respective Choose button.

For each investment round, this website randomly determines how much time you have to choose a company. In some rounds, you might have as little as 10 seconds, whereas for others you may take up to 45 seconds.

Every investment round ends as soon as you choose a company, or if you run out of time. The software will then calculate the total value of the company you have chosen, if any, and display the result.

Remember, at the very end of the study, we will pick one investment round at random, and you will be paid the value of the company that you chose in that round. If you did not choose a company in that round, then you will not receive a bonus payment.

Practice Round

We will next guide you through one practice round without a time limit, so that you can familiarize yourself with the interface.

Your performance-based reward is independent of what you do in this practice round.

Practice Round — Research Stage

Calculate each company's profit.

Remember: Profit = (Stores × Revenue per Store) - Total Expenses



Calculate and memorize the profit of each of the companies above. You will encounter the same three companies again in the choice stage. Any incorrect entries must be corrected before you will be able to proceed to the next stage.

Practice Round — Wait

You'll proceed to the choice stage in ____ seconds ...

While you wait, try to remember the profits of the three companies that you encountered during the research stage.

Practice Round - Choice Stage

Time Left: [Note clock]

Choose a company for your investment.

Remember: Value = Profit × Growth Factor

Company Name
Stores: ___
Revenue per Store: ___
Total Expenses: ___
Growth Factor: ?
Evaluate

Company Name
Stores: ___
Revenue per Store: ___
Total Expenses: ___
Growth Factor: ?
Evaluate

Company Name
Stores: ___
Revenue per Store: ___
Total Expenses: ___
Growth Factor: ?
Evaluate

Company Name
Stores: ___
Revenue per Store: ___
Total Expenses: ___
Growth Factor: ?
Evaluate

Company Name
Stores: ___
Revenue per Store: ___
Total Expenses: ___
Growth Factor: ?
Evaluate

Company Name
Stores: ___
Revenue per Store: ___
Total Expenses: ___
Growth Factor: ?
Evaluate

You must now choose the company in which you want to invest. Before selecting a company, however, you need to reveal its growth factor by clicking on the Evaluate button. You may evaluate as many companies as you'd like before making a final choice.

Your Choice

Company Name
Stores: ___
Revenue per Store: ___
Total Expenses: ___
Growth Factor: ___

Total Value: ___

This screen shows you the total value of the company that you have chosen. Thus, if you were to be paid based on your choice in this round, you would have earned a performance bonus of ____ USD.

Ready?

This concludes the practice round. By clicking on the Next button below, you will proceed to the first of 20 investment rounds.

In each investment round, you will encounter three unfamiliar companies during the research stage. In the corresponding choice stage, you will encounter the same three companies plus three new ones. Whether a particular company appears in both stages or only in the choice stage is entirely random.

For each company, we've chosen the number of stores, revenue per store, and total expenses randomly, subject to the constraint that profit is a round number between 2 and 10. Similarly, the software randomly picks a growth factor between 0.5 and 1.5 for every company.

You can thus expect to see any value of profits between 2 and 10, and any growth factor between 0.5 and 1.5 with equal probability.

Investment Round 1 of 20 — Research Stage

Calculate each company's profit.
Remember: Profit = (Stores × Revenue per Store) - Total Expenses



Please Wait

You'll proceed to the choice stage in ____ seconds ...

Investment Round 1 of 20 — Choice Stage

Time Left: ____

Choose a company for your investment.

Remember: Value = Profit × Growth Factor

Company Name
Stores: ___
Revenue per Store: ___
Total Expenses: ___
Growth Factor: ?
Evaluate

Company Name
Stores: ___
Revenue per Store: ___
Total Expenses: ___
Growth Factor: ?
Evaluate

Company Name
Stores: ___
Revenue per Store: ___
Total Expenses: ___
Growth Factor: ?
Evaluate

Company Name
Stores: ___
Revenue per Store: ___
Total Expenses: ___
Growth Factor: ?
Evaluate

Company Name
Stores: ___
Revenue per Store: ___
Total Expenses: ___
Growth Factor: ?
Evaluate

Company Name
Stores: ___
Revenue per Store: ___
Total Expenses: ___
Growth Factor: ?
Evaluate

Your Choice

Company Name
Stores: ___
Revenue per Store: ___
Total Expenses: ___
Growth Factor: ___

Total Value: ___

⋮ ⋮ ⋮

⋮ ⋮ ⋮

Results

Round	Your Choice	Stores	Revenue	Expenses	Growth	Total Value
1	Navy Blue	5	5	22	0.72	2.16
⋮	⋮	⋮	⋮	⋮	⋮	⋮
⋮	⋮	⋮	⋮	⋮	⋮	⋮
⋮	⋮	⋮	⋮	⋮	⋮	⋮
⋮	⋮	⋮	⋮	⋮	⋮	⋮
20	Spring Green	1	7	5	0.57	1.14

You will be paid based on round ____. You have, therefore, earned a performance-based bonus of ____ USD. This bonus is in addition to your base payment.

Thank You

Thank you for contributing to our research. Before you go, we have just one more question that will help us improve this study.

Did you feel like the instructions you received were clear?

- Yes
- No

If you have any other comments, please enter them below. We would be very interested in hearing your feedback.

Goodbye!

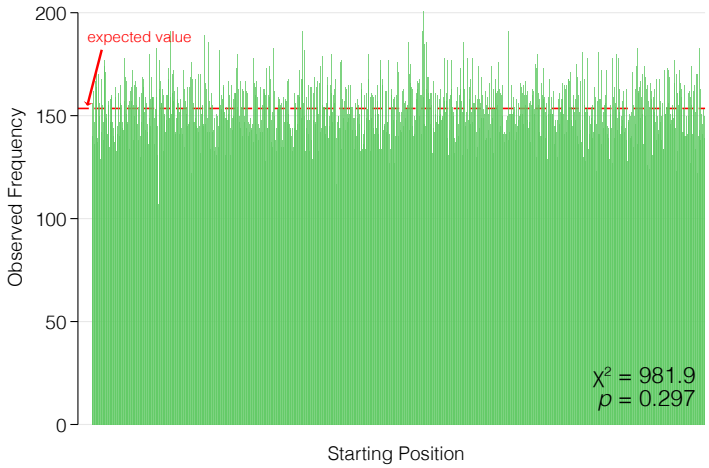
Thanks again for participating in our study!

Your Prolific completion code is: ____

Please click on the button below to be redirected to Prolific.

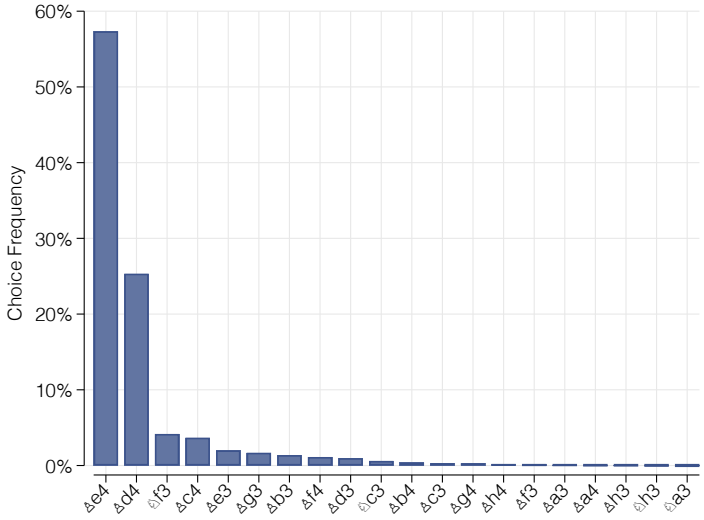
Appendix Figures

Appendix Figure AF.1: Distribution of Starting Positions in Chess960



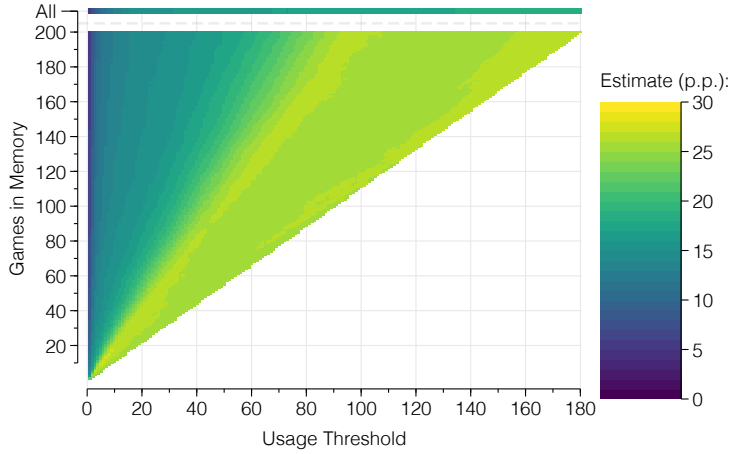
Notes: Figure shows the frequency of each starting position in our main sample of players' first game of Chess960. The reported χ^2 test refers to the null hypothesis of no differences across starting positions.

Appendix Figure AF.2: Choice Frequency of Opening Moves in Standard Chess



Notes: Figure shows the frequency with which the players in our main sample chose each opening move in prior games of standard chess.

Appendix Figure AF.3: Estimated Memory Premia, Controlling for the Set of Other In-Memory Moves

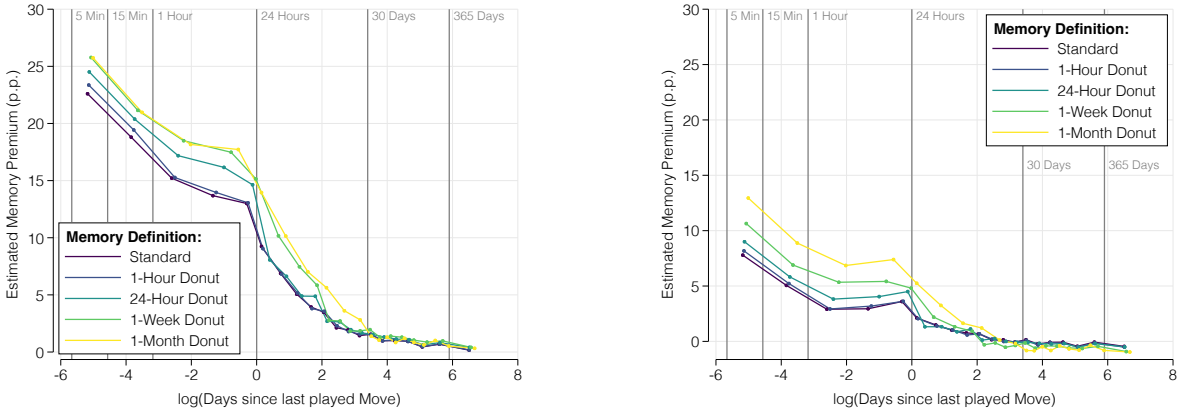


Notes: See Figure 7 in the main text. The only difference between this figure and that in the text is that the results above do not only control for the number of other in-memory moves but for *which* other moves are in-memory. More specifically, the estimates above are based on a regression model that replaces the move-by-starting-position-by-size-of-memory fixed effect in the specification in col. (6) of Table 2 with a move-by-starting-position-by-set-of-other-in-memory-moves fixed effect.

Appendix Figure AF.4: Decay in the Memory Premium, Based on Donut Definitions of Memory

(a) Baseline Set of Controls

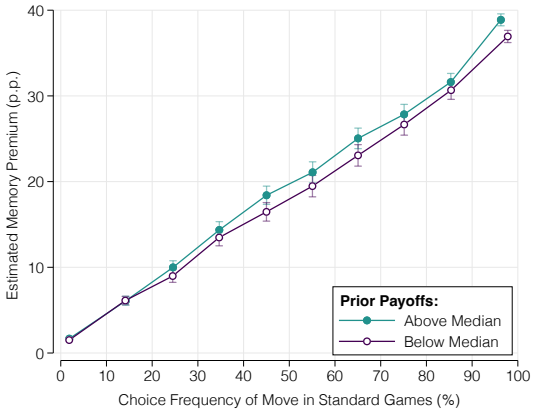
(b) Baseline and Frequency Controls



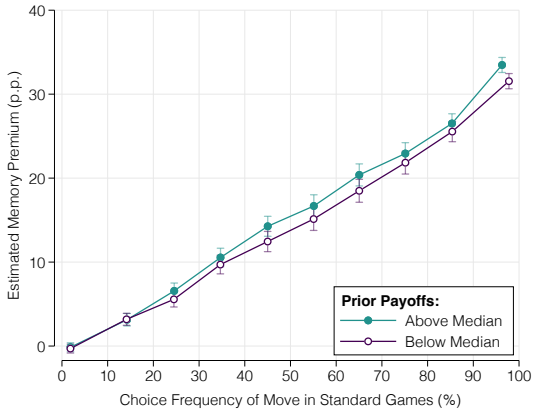
Notes: Figure replicates the analysis in Figure 9(b) in the text based on "donut definitions" of memory. According to these definitions, a move is defined as in-memory if, and only if, it had been played in at least one standard chess game on Lichess that occurred sufficiently long prior to a player's first game of Chess960. Panel (a) shows estimates of the memory premium (according to these alternative definitions) based on the regression model in col. (6) of Table 2. The estimates in panel (b) additionally control for the frequency with which a given move has previously been used. Error bars correspond to 95%-confidence intervals, accounting for clustering by player.

Appendix Figure AF.5: Frequency Effects, By Prior Success

(a) Baseline Set of Controls

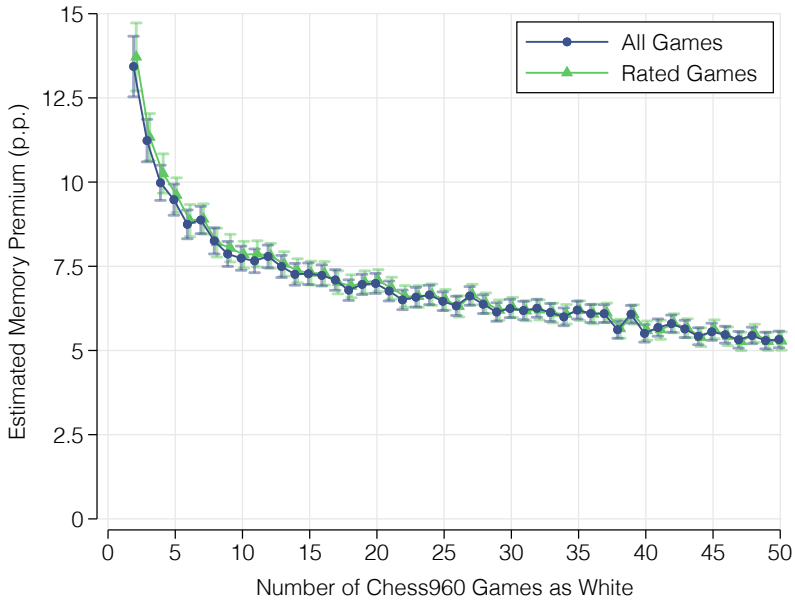


(b) Baseline and Frequency Controls



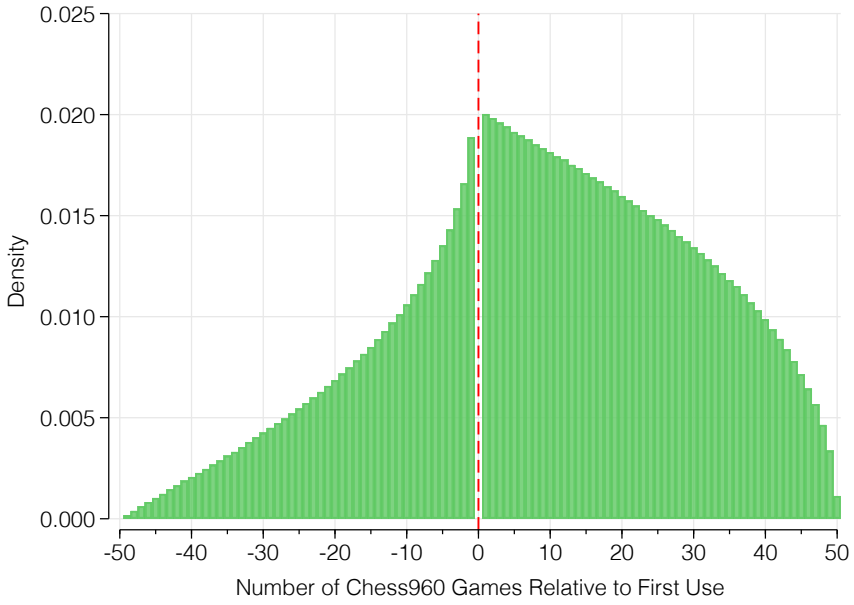
Notes: Figure replicates the analysis in Figure 9(a) in the text, differentiating between moves with above- and below-median player-specific average payoffs in prior use in standard chess games on Lichess. Payoffs are calculated by averaging over the outcomes of all games in which a player previously chose the respective move. Wins are valued at one point, draws at half a point, and losses at zero. Panel (a) shows estimates of the memory premium based on the regression model in col. (6) of Table 2. The estimates in panel (b) additionally control for how long ago the respective move has last been used in standard chess. Error bars correspond to 95%-confidence intervals, accounting for clustering by player.

Appendix Figure AF.6: Dynamics of the Memory Premium, Based on Prior Moves in Chess960



Notes: This figure replicates Figure 12 in the text for a definition of memory that is based on players' choices in previous games of Chess960.

Appendix Figure AF.7: Density of Observations around RD Threshold



Notes: Figure shows the density of observations around the threshold in the regression discontinuity design in Section 6.

Appendix Tables

Appendix Table AT.1: Memory Premium, Controlling for All Other In-Memory Moves

	Probability of Choosing Move					
	(1)	(2)	(3)	(4)	(5)	(6)
In-Memory (÷ 100)	12.07	3.83	3.83	4.94	4.96	5.03
	(0.04)	(0.04)	(0.04)	(0.13)	(0.13)	(0.13)
Centipawn Rating (÷ 100)		0.03				
		(0.00)				
Fixed Effects:						
Move	No	Yes	No	No	No	No
Move × Position	No	No	Yes	No	No	No
Move × Position × Set of Other In-Memory Moves	No	No	No	Yes	Yes	Yes
Player × Opponent Strength	No	No	No	No	Yes	Yes
Number of Standard Games as White	No	No	No	No	No	Yes
Mean of LHS Variable	5.08	5.08	5.08	5.08	5.08	5.08
R^2	0.05	0.13	0.16	0.67	0.67	0.67
N	2,898,455	2,898,455	2,898,455	2,898,455	2,898,455	2,898,455

Notes: See Table 2 in the text. The only difference between this table and Table 2 is that cols. (4)–(6) control for the set of other in-memory moves rather than the size of this set.

Appendix Table AT.2: Memory Premium, by Quality of Move in a Particular Starting Position

		Estimated Memory Premium		
Move	Choice Frequency	All Starting Positions	Close to Optimal	Suboptimal
♙a3	0.3	0.97	0.87	1.01
		(0.14)	(0.61)	(0.15)
♙a4	0.5	1.56	2.15	1.42
		(0.15)	(0.37)	(0.17)
♘a3	0.1	1.06	0.30	1.21
		(0.44)	(1.30)	(0.46)
♘c3	5.7	3.51	3.76	3.27
		(0.42)	(0.65)	(0.57)
♙b3	5.5	5.02	7.57	3.68
		(0.21)	(0.48)	(0.21)
♙b4	2.0	1.25	1.52	1.16
		(0.18)	(0.47)	(0.19)
♙c3	1.9	1.13	1.60	1.02
		(0.16)	(0.44)	(0.16)
♙c4	4.5	3.41	3.90	2.92
		(0.15)	(0.26)	(0.19)
♙d3	2.5	2.12	1.84	2.18
		(0.12)	(0.29)	(0.14)
♙d4	17.4	11.64	12.48	10.47
		(0.23)	(0.32)	(0.34)
♙e3	3.3	2.73	2.76	2.75
		(0.11)	(0.23)	(0.13)
♙e4	35.1	29.56	28.49	31.23
		(0.39)	(0.55)	(0.54)
♙f3	1.9	1.04	1.09	1.05
		(0.15)	(0.46)	(0.16)
♙f4	4.3	1.87	1.58	2.04
		(0.16)	(0.26)	(0.20)
♘f3	8.8	7.81	8.17	7.45
		(0.40)	(0.61)	(0.53)
♘h3	0.1	2.09	–1.08	2.16
		(0.42)	(0.65)	(0.46)
♙g3	6.3	5.25	7.08	4.39
		(0.20)	(0.46)	(0.21)
♙g4	2.4	1.49	1.28	1.55
		(0.18)	(0.49)	(0.18)
♙h3	0.4	1.39	2.02	1.30
		(0.14)	(0.80)	(0.15)
♙h4	0.7	2.87	3.09	2.76
		(0.17)	(0.49)	(0.18)

Notes: Entries are average choice frequencies and estimated memoria premia for all standard opening moves, either in all starting positions or in starting positions in which the move is and is not close to optimal (i.e., within 20 centipawns of the optimal move in a given position). All estimates are based on the regression model and sample in col. (6) of Table 2, allowing for heterogeneity in the memory premium across opening moves. Estimates are scaled to correspond to the percentage-point change in choice probability associated with a one-unit increase in the respective regressor. Standard errors are clustered by player and shown in parentheses.

Appendix Table AT.3: Memory Premium, Flipped Standard Board

	Probability of Choosing Move					
	(1)	(2)	(3)	(4)	(5)	(6)
Move In-Memory (÷ 100)	3.13	2.93	3.76	3.63	3.73	3.62
	(1.04)	(1.29)	(1.58)	(1.62)	(1.60)	(1.63)
Mirror Move In-Memory (÷ 100)		0.65		1.87		1.87
		(1.24)		(1.52)		(1.52)
Fixed Effects:						
Move	Yes	Yes	No	No	No	No
Move × Size of Memory	No	No	Yes	Yes	Yes	Yes
Difference in Players' Strength	No	No	No	No	Yes	Yes
Previous Standard Games as White	No	No	No	No	Yes	Yes
Mean of LHS Variable	5.00	5.00	5.00	5.00	5.00	5.00
R^2	0.30	0.30	0.39	0.39	0.39	0.39
N	3,340	3,340	3,340	3,340	3,340	3,340

Notes: Entries are coefficients and standard errors from estimating variants of the regression model in eq. (1) by ordinary least squares. The sample consists of all available opening moves in players' first game of Chess960 as White, restricting attention to players who are shown the mirrored standard board (i.e., starting position 534). All estimates are scaled to correspond to the percentage-point change in choice probability associated with a one-unit increase in the respective regressor. Standard errors are clustered by player, and are shown in parentheses.

Appendix Table AT.4: Memory Premium, Boards with Corner Royalty

A. All Moves

	Probability of Choosing Move					
	(1A)	(2A)	(3A)	(4A)	(5A)	(6A)
In-Memory (÷ 100)	11.12	3.91	3.89	4.63	4.62	4.71
	(0.15)	(0.13)	(0.13)	(0.17)	(0.18)	(0.18)
Centipawn Rating (÷ 100)		0.03				
		(0.00)				
Fixed Effects:						
Move	No	Yes	No	No	No	No
Move × Position	No	No	Yes	No	No	No
Move × Position × Size of Memory	No	No	No	Yes	Yes	Yes
Player × Opponent Strength	No	No	No	No	Yes	Yes
Previous Standard Games as White	No	No	No	No	No	Yes
Mean of LHS Variable	5.06	5.06	5.06	5.06	5.06	5.06
R^2	0.04	0.11	0.13	0.21	0.21	0.21
N	219,952	219,952	219,952	219,952	219,952	219,952

B. Centerpawns Only

	Probability of Choosing Move					
	(1B)	(2B)	(3B)	(4B)	(5B)	(6B)
In-Memory (÷ 100)	16.71	6.81	6.78	7.76	7.37	6.97
	(0.26)	(0.26)	(0.26)	(0.32)	(0.40)	(0.44)
Centipawn Rating (÷ 100)		0.04				
		(0.01)				
Fixed Effects:						
Move	No	Yes	No	No	No	No
Move × Position	No	No	Yes	No	No	No
Move × Position × Size of Memory	No	No	No	Yes	Yes	Yes
Player × Opponent Strength	No	No	No	No	Yes	Yes
Previous Standard Games as White	No	No	No	No	No	Yes
Mean of LHS Variable	13.42	13.42	13.42	13.42	13.42	13.42
R^2	0.06	0.14	0.15	0.23	0.31	0.33
N	44,488	44,488	44,488	44,488	44,488	44,488

Notes: Entries are coefficients and standard errors from estimating variants of the regression model in eq. (1) by ordinary least squares. In Panel A, the sample consists of all opening moves in starting postions in which the king and queen sit on squares a1, b1, g1, or h1. In Panel B, the sample is restricted to opening moves ♙d3, ♙d4, ♙e3, and ♙e4 and to positions in which the king and queen sit on squares a1, b1, g1, or h1. All estimates are scaled to correspond to the percentage-point change in choice probability associated with a one-unit increase in the respective regressor. Standard errors are clustered by player, and are shown in parentheses.

Appendix Table AT.5: Memory Premium, 1-Hour Donut-Definition of Memory

	Probability of Choosing Move					
	(1)	(2)	(3)	(4)	(5)	(6)
In-Memory (÷ 100)	12.09	3.82	3.82	4.51	4.51	4.52
	(0.04)	(0.04)	(0.04)	(0.05)	(0.05)	(0.05)
Centipawn Rating (÷ 100)		0.03				
		(0.00)				
Fixed Effects:						
Move	No	Yes	No	No	No	No
Move × Position	No	No	Yes	No	No	No
Move × Position × Size of Memory	No	No	No	Yes	Yes	Yes
Player × Opponent Strength	No	No	No	No	Yes	Yes
Previous Standard Games as White	No	No	No	No	No	Yes
Mean of LHS Variable	5.08	5.08	5.08	5.08	5.08	5.08
R^2	0.05	0.13	0.16	0.24	0.24	0.24
N	2,898,455	2,898,455	2,898,455	2,898,455	2,898,455	2,898,455

Notes: See Table 2 in the text. The only difference between this table and Table 2 is that we apply a donut definition of memory, according to which a move is said to be in-memory if, and only if, it had been played in at least one standard chess game on Lichess that occured at least one hour prior to a player's first game of Chess960.

Appendix Table AT.6: Memory Premium, 1-Day Donut-Definition of Memory

	Probability of Choosing Move					
	(1)	(2)	(3)	(4)	(5)	(6)
In-Memory (÷ 100)	12.15	3.83	3.82	4.75	4.75	4.75
	(0.04)	(0.04)	(0.04)	(0.05)	(0.05)	(0.05)
Centipawn Rating (÷ 100)		0.03				
		(0.00)				
Fixed Effects:						
Move	No	Yes	No	No	No	No
Move × Position	No	No	Yes	No	No	No
Move × Position × Size of Memory	No	No	No	Yes	Yes	Yes
Player × Opponent Strength	No	No	No	No	Yes	Yes
Previous Standard Games as White	No	No	No	No	No	Yes
Mean of LHS Variable	5.08	5.08	5.08	5.08	5.08	5.08
R^2	0.05	0.13	0.16	0.24	0.24	0.24
N	2,898,455	2,898,455	2,898,455	2,898,455	2,898,455	2,898,455

Notes: See Table 2 in the text. The only difference between this table and Table 2 is that we apply a donut definition of memory, according to which a move is said to be in-memory if, and only if, it had been played in at least one standard chess game on Lichess that occured at least one day prior to a player's first game of Chess960.

Appendix Table AT.7: Memory Premium, 1-Week Donut-Definition of Memory

	Probability of Choosing Move					
	(1)	(2)	(3)	(4)	(5)	(6)
In-Memory (÷ 100)	11.94	3.54	3.53	5.16	5.17	5.17
	(0.04)	(0.04)	(0.04)	(0.06)	(0.06)	(0.06)
Centipawn Rating (÷ 100)		0.03				
		(0.00)				
Fixed Effects:						
Move	No	Yes	No	No	No	No
Move × Position	No	No	Yes	No	No	No
Move × Position × Size of Memory	No	No	No	Yes	Yes	Yes
Player × Opponent Strength	No	No	No	No	Yes	Yes
Previous Standard Games as White	No	No	No	No	No	Yes
Mean of LHS Variable	5.08	5.08	5.08	5.08	5.08	5.08
R^2	0.05	0.13	0.16	0.24	0.24	0.24
N	2,898,455	2,898,455	2,898,455	2,898,455	2,898,455	2,898,455

Notes: See Table 2 in the text. The only difference between this table and Table 2 is that we apply a donut definition of memory, according to which a move is said to be in-memory if, and only if, it had been played in at least one standard chess game on Lichess that occured at least one week prior to a player's first game of Chess960.

Appendix Table AT.8: Memory Premium, 1-Month Donut-Definition of Memory

	Probability of Choosing Move					
	(1)	(2)	(3)	(4)	(5)	(6)
In-Memory (÷ 100)	11.29	2.87	2.87	5.16	5.17	5.19
	(0.05)	(0.04)	(0.04)	(0.07)	(0.07)	(0.07)
Centipawn Rating (÷ 100)		0.03				
		(0.00)				
Fixed Effects:						
Move	No	Yes	No	No	No	No
Move × Position	No	No	Yes	No	No	No
Move × Position × Size of Memory	No	No	No	Yes	Yes	Yes
Player × Opponent Strength	No	No	No	No	Yes	Yes
Previous Standard Games as White	No	No	No	No	No	Yes
Mean of LHS Variable	5.08	5.08	5.08	5.08	5.08	5.08
R^2	0.03	0.13	0.16	0.23	0.23	0.23
N	2,898,455	2,898,455	2,898,455	2,898,455	2,898,455	2,898,455

Notes: See Table 2 in the text. The only difference between this table and Table 2 is that we apply a donut definition of memory, according to which a move is said to be in-memory if, and only if, it had been played in at least one standard chess game on Lichess that occured at least one month prior to a player's first game of Chess960.

Appendix Table AT.9: Summary Statistics for Experimental Data

			Percentile					
Variable	Mean	SD	5%	25%	50%	75%	95%	*N*
A. Subject Characteristics								
Age	41.0	13.4	22.0	31.0	38.0	50.0	66.0	501
Female (%)	48.9							501
Instructions Clear (%)	98.4							501
B. Choice Variables across Players								
Alternatives Evaluated (%)	73.0	29.3	16.7	48.3	85.0	98.2	100.0	501
Share In-Memory among Chosen Alternatives (%)	59.9	17.3	35.0	47.4	57.9	70.0	94.7	501
Share In-Memory among Evaluated Alternatives (%)	54.6	11.5	46.2	50.0	50.4	54.8	85.4	501
Share In-Memory among Alternatives Evaluated First (%)	65.7	20.9	36.8	50.0	63.2	84.2	100.0	501
C. Outcomes across Players								
Average Value of Chosen Alternatives	8.6	1.2	6.6	7.9	8.7	9.5	10.4	501
Average Value of Evaluated Alternatives	6.4	0.8	5.6	5.9	6.1	6.5	8.2	501
Average Value of Alternatives Evaluated First	6.9	1.2	5.1	6.0	6.8	7.7	8.9	501

Notes: This table displays summary statistics for selected variables in the data from our experiment. All observations are at the participant level. Each observation in Panel B and C corresponds to a participant-level average across all 20 rounds of the choice task.

Appendix Table AT.10: Experimental Results, Including Choice Failures

A. Main Sample

	Probability Chosen				Probability Evaluated				Probability Evaluated First			
	(1A)	(2A)	(3A)	(4A)	(5A)	(6A)	(7A)	(8A)	(9A)	(10A)	(11A)	(12A)
In-Memory (÷ 100)	6.60	6.82	6.81	6.81	7.06	7.10	7.10	7.10	10.58	10.73	10.74	10.74
	(0.51)	(0.52)	(0.52)	(0.52)	(0.64)	(0.64)	(0.64)	(0.64)	(0.63)	(0.63)	(0.62)	(0.62)
Fixed Effects:												
Own × Other Profit	No	Yes	Yes	Yes	No	Yes	Yes	Yes	No	Yes	Yes	Yes
Display Position	No	No	Yes	Yes	No	No	Yes	Yes	No	No	Yes	Yes
Time Limit	No	No	No	Yes	No	No	No	Yes	No	No	No	Yes
Mean of LHS Variable	16.67	16.67	16.67	16.67	73.25	73.25	73.25	73.25	16.67	16.67	16.67	16.67
R^2	0.01	0.22	0.22	0.22	0.01	0.23	0.24	0.25	0.02	0.20	0.39	0.39
N	57,792	57,792	57,792	57,792	57,792	57,792	57,792	57,792	57,792	57,792	57,792	57,792

B. Incl. Choice Failures

	Probability Chosen				Probability Evaluated				Probability Evaluated First			
	(1B)	(2B)	(3B)	(4B)	(5B)	(6B)	(7B)	(8B)	(9B)	(10B)	(11B)	(12B)
In-Memory (÷ 100)	6.34	6.55	6.55	6.55	6.94	7.02	7.01	7.01	10.42	10.59	10.54	10.54
	(0.49)	(0.50)	(0.50)	(0.50)	(0.63)	(0.63)	(0.63)	(0.63)	(0.61)	(0.62)	(0.60)	(0.60)
Fixed Effects:												
Own × Other Profit	No	Yes	Yes	Yes	No	Yes	Yes	Yes	No	Yes	Yes	Yes
Display Position	No	No	Yes	Yes	No	No	Yes	Yes	No	No	Yes	Yes
Time Limit	No	No	No	Yes	No	No	No	Yes	No	No	No	Yes
Mean of LHS Variable	16.02	16.02	16.02	16.02	72.48	72.48	72.48	72.48	16.52	16.52	16.52	16.52
R^2	0.01	0.22	0.22	0.22	0.01	0.23	0.23	0.24	0.02	0.19	0.38	0.38
N	60,120	60,120	60,120	60,120	60,120	60,120	60,120	60,120	60,120	60,120	60,120	60,120

Notes: Entries are coefficients and standard errors from ordinary least squares regressions. In columns 1–4, the outcome is an indicator for whether an alternative is chosen. The outcome in columns 5–8 is an indicator for whether an alternative is evaluated, while that in columns 9–12 is an indicator for whether the alternative is evaluated first. Following the pre-registration, panel A includes all alternatives from choice tasks in which participants submitted a timely choice. Panel B includes all observations, including observations from choice tasks in which participants ran out of time. In that case, none of the available alternatives are coded as having been chosen. All estimates are scaled to correspond to the percentage-point change in the probability of the corresponding outcome associated with a one-unit increase in the respective regressor. Standard errors are clustered at the participant level, and are shown in parentheses.

Appendix Table AT.11: Experimental Results, by Ranking of Alternatives in the Research Stage

	Probability Chosen				Probability Evaluated				Probability Evaluated First			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Best In-Memory (÷ 100)	22.04	16.37	16.37	16.37	17.21	14.67	14.67	14.71	25.68	24.54	24.74	24.74
	(1.01)	(1.05)	(1.05)	(1.05)	(1.15)	(1.09)	(1.09)	(1.09)	(1.34)	(1.34)	(1.31)	(1.31)
Second Best In-Memory (÷ 100)	–1.16	1.78	1.79	1.79	2.54	3.70	3.74	3.71	1.49	2.54	2.79	2.79
	(0.49)	(0.51)	(0.51)	(0.51)	(0.67)	(0.69)	(0.69)	(0.69)	(0.47)	(0.53)	(0.48)	(0.48)
Third Best In-Memory (÷ 100)	–6.07	–0.46	–0.48	–0.48	–1.96	0.63	0.58	0.58	–0.04	1.28	0.75	0.75
	(0.43)	(0.45)	(0.45)	(0.45)	(0.56)	(0.55)	(0.55)	(0.55)	(0.48)	(0.54)	(0.45)	(0.45)
Fixed Effects:												
Own × Other Profit	No	Yes	Yes	Yes	No	Yes	Yes	Yes	No	Yes	Yes	Yes
Display Position	No	No	Yes	Yes	No	No	Yes	Yes	No	No	Yes	Yes
Time Limit	No	No	No	Yes	No	No	No	Yes	No	No	No	Yes
Mean of LHS Variable	16.67	16.67	16.67	16.67	73.25	73.25	73.25	73.25	16.67	16.67	16.67	16.67
R^2	0.06	0.23	0.23	0.23	0.02	0.24	0.24	0.25	0.07	0.22	0.41	0.41
N	57,792	57,792	57,792	57,792	57,792	57,792	57,792	57,792	57,792	57,792	57,792	57,792

Notes: Entries are coefficients and standard errors from ordinary least squares regressions. In columns 1–4, the outcome is an indicator for whether an alternative is chosen. The outcome in columns 5–8 is an indicator for whether an alternative is evaluated, while that in columns 9–12 is an indicator for whether the alternative is evaluated first. The regressors correspond to indicators for whether the alternative has the highest, second highest, or third highest expected value among all in-memory options. The sample includes all alternatives from choice tasks in which participants submitted a timely choice. All estimates are scaled to correspond to the percentage-point change in the probability of the corresponding outcome associated with a one-unit increase in the respective regressor. Standard errors are clustered at the participant level, and are shown in parentheses.

Appendix Table AT.12: Results from Robustness Experiment

	Outcome		
	Chosen	Evaluated	Evaluated First
A. All In-Memory Alternatives			
In-Memory	2.55	2.71	5.67
	(0.75)	(0.79)	(0.82)
B. Most Promising In-Memory Alternative			
In-Memory × Expected Value = 10	27.94	9.57	15.94
	(1.80)	(1.65)	(2.08)
In-Memory × Expected Value = 9	22.82	10.43	15.70
	(1.82)	(1.59)	(2.06)
In-Memory × Expected Value = 8	16.52	7.70	15.56
	(2.04)	(1.59)	(2.24)
In-Memory × Expected Value = 7	9.73	8.99	10.63
	(2.21)	(1.60)	(2.07)
In-Memory × Expected Value = 6	3.82	1.72	7.23
	(2.09)	(1.77)	(2.30)
In-Memory × Expected Value = 5	7.66	3.63	5.87
	(2.86)	(2.62)	(2.67)
In-Memory × Expected Value = 4	2.62	7.08	4.80
	(3.19)	(2.65)	(3.21)
In-Memory × Expected Value = 3	–5.25	2.11	5.01
	(3.33)	(5.24)	(4.16)
In-Memory × Expected Value = 2	–10.63	–8.66	0.46
	(4.54)	(13.89)	(7.43)
C. Other In-Memory Alternatives			
In-Memory × Expected Value = 9	8.67	7.63	6.47
	(2.54)	(1.93)	(2.58)
In-Memory × Expected Value = 8	2.74	2.94	2.69
	(1.80)	(1.61)	(1.54)
In-Memory × Expected Value = 7	–1.77	1.66	0.61
	(1.38)	(1.40)	(1.45)
In-Memory × Expected Value = 6	–5.42	3.32	1.61
	(1.26)	(1.11)	(1.38)
In-Memory × Expected Value = 5	–8.17	–0.98	0.80
	(1.01)	(1.29)	(1.27)
In-Memory × Expected Value = 4	–10.32	–3.35	0.44
	(0.86)	(1.43)	(1.20)
In-Memory × Expected Value = 3	–11.38	–2.79	–0.30
	(0.83)	(1.35)	(1.12)
In-Memory × Expected Value = 2	–10.75	–3.36	0.33
	(0.85)	(1.45)	(1.07)

Notes: Entries are coefficients and standard errors from ordinary least squares regressions replicating the results in panels A and B of Figure 15 based on the data from the robustness experiment described in Appendix E. The estimates thus correspond to the percentage point difference in the respective outcome between particular in-memory options and the average out-of-memory alternative. The sample corresponds to all alternatives from choice tasks in which participants submitted a timely choice. Standard errors are clustered at the participant level, and are shown in parentheses.

Appendix Table AT.13: Ancillary Experimental Results

	Log Deliberation Time						Number of Evaluations			Optimization Error		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Choose In-Memory Option	–0.13	–0.09	–0.08				–0.53	–0.24	–0.22			
	(0.02)	(0.01)	(0.01)				(0.07)	(0.03)	(0.03)			
Number of Evaluations				0.11	0.15	0.12				9.33	9.28	9.40
				(0.01)	(0.01)	(0.01)				(0.27)	(0.40)	(0.40)
Log Deliberation Time										–8.04	–3.51	1.32
										(1.08)	(1.14)	(1.23)
Fixed Effects:												
Subject	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
Time Limit	No	No	Yes	No	No	Yes	No	No	Yes	No	No	Yes
Mean of LHS Variable	2.56	2.56	2.56	2.56	2.56	2.56	4.39	4.39	4.39	32.40	32.40	32.40
R^2	0.01	0.46	0.57	0.18	0.55	0.62	0.02	0.69	0.71	0.15	0.27	0.28
N	9,632	9,632	9,632	9,632	9,632	9,632	9,632	9,632	9,632	9,632	9,632	9,632

Notes: Entries are coefficients and standard errors from ordinary least squares regressions. In columns 1–6, the outcome of interest is the logarithm of the time that a subject spent in the choice stage of a given round. The outcome in columns 7–9 is the number of evaluated alternatives in a given round, while that in columns 10–12 is an indicator for choosing an alternative that is suboptimal relative to the best option that the subject evaluated. The regressors shown on the left correspond to an indicator for whether an alternative in in-memory on account of having been encountered in the research stage, the logarithm of the deliberation time, and the number of evaluated alternatives. The estimates in columns 10–12 are scaled to correspond to the percentage-point change in the probability of choosing a suboptimal alternative associated with a one-unit increase in the respective regressor. The sample includes all alternatives from choice tasks in which participants submitted a timely choice. Standard errors are clustered at the participant level, and are shown in parentheses.