NBER WORKING PAPER SERIES

COMPLEXITY AND CHOICE

Yuval Salant Jorg L. Spenkuch

Working Paper 30002 http://www.nber.org/papers/w30002

NATIONAL BUREAU OF ECONOMIC RESEARCH 1050 Massachusetts Avenue Cambridge, MA 02138 April 2022

We are grateful to Sandeep Baliga, Tim Feddersen, Peter Klibanoff, Daniel Martin, Pablo Montagnes, and David Smerdon, as well as audiences at Boston College, Monash University, Northwestern, University of Alicante, UC Santa Barbara, and SITE for helpful suggestions. This research was supported in part through the computational resources provided by the Quest high-performance computing facility at Northwestern University. 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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2022 by Yuval Salant and Jorg L. Spenkuch. 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.

Complexity and Choice Yuval Salant and Jorg L. Spenkuch NBER Working Paper No. 30002 April 2022 JEL No. D00,D01,D03,D9,D90

ABSTRACT

We develop a model of satisficing with evaluation errors that incorporates complexity at the level of individual alternatives. We test the model predictions in a novel data set with information on hundreds of millions of chess moves by experienced players. Consistent with the theory, complex optimal moves are chosen less frequently than simpler ones. Choice frequencies of suboptimal moves follow the opposite pattern. The former finding distinguishes satisficing from a large class of maximization-based models. We further document that skill and time moderate the adverse effect of complexity, and that they complement each other in doing so. Finally, we provide evidence that suboptimal behavior also hinges on the composition of the choice set but not its size. Our findings help to shed some of the first light on the importance of complexity outside of the laboratory.

Yuval Salant Kellogg School of Management 2211 Campus Dr Evanston, IL 60208 y-salant@kellogg.northwestern.edu

Jorg L. Spenkuch Kellogg School of Management Northwestern University 2211 Campus Drive Evanston, IL 60208 and NBER j-spenkuch@kellogg.northwestern.edu

A data appendix is available at http://www.nber.org/data-appendix/w30002

1. Introduction

In any given choice situation some of the available alternatives may be inherently complex and, therefore, hard to evaluate. Insurance contracts, for instance, might consist of tens or even hundreds of clauses, which collectively make it difficult to assess a particular plan. Durable goods might have dozens of attributes that jointly determine value, and policy proposals often contain a great number of details that bedevil decision-making. Strategies in dynamic games can include so many contingencies that identifying an optimal one becomes nearly impossible. The goal of this paper is to better understand how complexity of this kind affects choice behavior.

Drawing on the notion of Kolmogorov complexity in information theory, we conceive of an object's complexity as the length of its efficient representation. For example, in a dynamic game, the inherent complexity of an action may correspond to the size of the ensuing subgame. The complexity of a contract might be approximated by the number of non-redundant clauses; and in the context of durable goods, complexity may correspond to the length of the array of relevant attributes.¹ We propose that alternatives with longer efficient representations are more challenging to evaluate, and that this difficulty manifests itself in decision-making as errors in the perception of value. A decision maker may only be able to precisely evaluate alternatives with relatively short representations. When an alternative's representation exceeds a certain limit, evaluations become noisy.

Our analysis begins by incorporating this idea into an empirically testable theory of choice. The theory has two main ingredients. The first one adapts the familiar random-utility framework to study complexity. In the model, each alternative is characterized by its true value and its inherent object complexity, i.e., the length of its efficient representation. The decision maker (DM) relies on her skill to break down complexity and evaluate an alternative's value. By skill, we mean the DM's cognitive resources, experience, expertise, and so on. Skill formally corresponds to the maximal representation length the DM is able to evaluate without error. When an object's inherent complexity exceeds the DM's skill, she only obtains a noisy estimate of value. The accuracy of the estimate depends negatively on how much of the object's representation she was unable to process.

The second ingredient of our theory incorporates this noisy evaluation process into a model of choice. Here, we build on Simon's (1955; 1972) seminal work on bounded rationality and satisficing. According to Simon, individuals may not consider all possible options and pick

¹For longer object representations to imply higher complexity, all alternatives have to be described in the same language, although the language need not be the same across contexts. Unlike Kolmogorov complexity, we do not insist on the shortest possible representation. In fact, in some economic environments, strategic parties, such as firms, can deliberately simplify or complicate object descriptions in order to achieve their goals.

the best, but examine a rather small number, making a choice as soon as they find one alternative that they regard as satisfactory. In our model, the decision maker has in mind an aspiration level that she wishes to exceed. She lists all available alternatives in random order, and sequentially evaluates them until she encounters one whose estimated value exceeds the aspiration level. This is the alternative she chooses.

After developing the model's predictions, we test them in a novel dataset on chess endgames. Chess is an environment that admits an unambiguous measure of value as well as a natural measure of object complexity.

Every board configuration in a chess game corresponds to a choice set in which the alternatives are all available legal moves. By Zermelo's Theorem (1913), any chess move is of one of three types. A winning move allows the current player to move to force a win. A losing move enables her opponent to guarantee himself a win, whereas a drawing move lets both players force a draw. Although computing these types is generally infeasible in the opening and middlegame phases, endgames with up to six pieces have been definitively solved by modern computers. We can therefore measure the true value of all endgame moves.

As in any dynamic game, assessing the value of a chess move requires DMs to examine the ensuing subgame. Because larger game trees are likely harder to evaluate than smaller ones, we propose to equate the inherent complexity of a move with the size of the subsequent subgame. As an empirical proxy for the size of a subgame, we rely on what chess players refer to as depth-to-mate (DTM). DTM is a theoretical metric of how fast the dominant player can force a checkmate when the losing player resists as long as possible. In our data, moves' DTM is strongly correlated with the length of actual play, suggesting that DTM is a useful proxy for the size of the ensuing subgame and, therefore, the difficulty of evaluating the corresponding move.

Our data come from two sources. The first one is lichess.org, one of the three most popular internet chess servers. We have information on the universe of moves in all rated games on the platform from January 2013 through August 2020.² We augment these data with information from the Syzygy and Nalimov endgame tablebases (Nalimov et al. 2000; Man 2013). Both databases record the type of every available move in all non-trivial endgame positions with up to six pieces—just as in Zermelo's Theorem. The Nalimov database also contains information on moves' DTM, our proxy for object complexity. Our analysis focuses on the choices of nearly a quarter million experienced users, ranging from seasoned hobbyists to the best players in the world. In total, we observe about 227 million choices from sets with approximately 4.6 billion alternatives.

²Rated games are consequential in the sense that their outcomes directly affect users' strength ratings and rankings on the site. Anecdotal evidence suggests players care intensely about their ratings.

The first part of our empirical investigation examines how object complexity affects choice behavior. Comparing moves of the same type across otherwise similar choice sets, we observe that winning moves become less likely to be chosen as their DTM increases. For losing moves, the reverse holds. To see the economic significance of this result, consider a two-node increase in the DTM of a winning move. According to our estimates, such an increase corresponds to a 1.0–2.3 p.p. decrease in the choice frequency of the corresponding move, which translates to a 10%–24% reduction relative to the mean choice frequency of winning moves.

We further document that, within the same the choice set, winning moves with a low DTM are more often selected than their counterparts with higher DTM. The opposite holds for losing moves. We prove that the former finding is inconsistent with the predictions of maximization-based models in which the DM conducts noisy evaluations of all alternatives before choosing the one with the maximal estimated value. Our theory thus yields a new empirical test that leverages object complexity to distinguish satisficing from noisy maximization.

In the second part of the empirical analysis, we probe the model predictions about skill and complexity. Given that our data contain approximately 3.4 million moves by chess masters, of which nearly 350,000 were made by grandmasters, we can ask whether highly skilled players are less likely to make mistakes when object complexity is large. By mistake we mean choosing a drawing or losing move in the presence of at least one winning alternative. Consistent with the theory, we only detect large differences between titled and untitled players when the minimal DTM among winning moves is relatively high, i.e., when every move is hard to evaluate. Taking one step beyond the theory, we document that skill and time are complements in moderating the adverse effect of object complexity. When players are afforded more time to make decisions, titled players' frequency of suboptimal choices declines by more than that of untitled ones.

In the last part of the analysis, we examine set- rather than object-level factors that complicate decision making. Confirming previous work by Anderson et al. (2017), we document that chess players are more likely to make mistakes as the share of winning moves in the choice set decreases. Going beyond Anderson et al. (2017), we show that the *mix* of suboptimal alternatives is another key determinant of mistake probabilities. Consider, for example, replacing a single losing move in the choice set with a drawing alternative. Based on our estimates, such a change increases the frequency of mistakes on average by about .84 p.p.—or 14% relative to the mean frequency of mistakes. Contrary to psychological notions of choice overload (Iyengar and Lepper 2000), we find that what matters for the quality of decision making in this setting is the composition of the choice set, not its size. Once we carefully control for the fraction of winning, drawing, and losing moves in the set, the number of available alternatives and the frequency of mistakes are practically uncorrelated.

The work in this paper speaks directly to the theoretical literature on how complexity considerations affect outcomes in single- and multi-person environments. This research usually conceives of complexity as affecting behavior through constraints on agents' computational abilities and memory (e.g., Neyman 1985; Rubinstein 1986; Abreu and Rubinstein 1988; Kalai and Stanford 1988; Salant 2011; Wilson 2014; Jakobsen 2020). A high-level takeaway is that computational constraints can greatly affect both individual and strategic outcomes.

Our contribution relative to extant theoretical work is twofold. First, we draw on the notion of Kolmogorov complexity to develop a tractable model of choice that highlights objects' inherent complexity as a source of boundedly rational behavior. Second, we provide evidence on the empirical relevance of complexity considerations in general, and object complexity in particular.

In addition, our work complements a growing experimental literature on complexity in decision-making (see, e.g., Huck and Weizsäcker 1999; Gabaix et al. 2006; Abeler and Jäger 2015; Bossaerts and Murawski 2017; Enke and Graeber 2021). Rubinstein (2007, 2016), for instance, shows that decisions that require explicit cognitive reasoning take far longer to complete than those that do not. Caplin et al. (2011) provide evidence that individuals rely on satisficing in choice environments in which evaluating each option takes time and effort. Oprea (2020) develops a revealed-preference methodology to measure the cost of complexity. He finds subjects are willing to pay significant amounts in order to avoid tasks that are inherently complex. Overall, laboratory experiments support the idea that complexity can affect decision-making.

Outside of the laboratory, tests of fundamental decision-theoretic concepts remain rare.³ As Chiappori et al. (2002) note, nonexperimental settings may be intractable, with choice sets that need not be known in their entirety, or even be specified ex ante. Moreover, theoretical predictions may hinge on subtle properties of utility functions, intricacies of payoff structures, and individuals' beliefs—all of which are typically unobserved by the econometrician. As a result, we know little about how complexity affects decision-making in real-world environments.⁴

Chess endgames provide an almost ideal empirical setting to study this question. In addition to yielding observable variation in complexity and admitting an ordinal measure of alternatives'

³A related literature asks whether some of the basic tenets of game theory are consistent with observed behavior in different real-world environments. Walker and Wooders (2001), Chiappori et al. (2002), Palacios-Huerta (2003), and Hsu et al. (2007) all study minimax play in professional sports, while Spenkuch et al. (2018) examine backward induction in sequential voting. On the whole, the evidence from these settings corroborates theory more closely than one might have guessed based on an abundance of negative findings from the laboratory (see, e.g., Camerer 2003 for a review).

⁴Levitt and List (2007) argue that even if individuals consistently make mistakes in the laboratory, competitive forces and experience may limit such behavior in the real world.

value, chess possesses at least four additional attractive features. First, the rules of the game are known to players and there is virtually no uncertainty about primitives such as choice sets. Yet, evaluating different moves often strains the bounds of human cognition. Second, data on chess games are abundant. We therefore have enough statistical power to test even subtle theoretical predictions. Third, chess games are played under different time controls, and there are large differences in players' skill. This allows us to analyze how complexity, skill, and time interact. Fourth, we study experienced players in a familiar environment, thus minimizing the risk that our findings are due to an unfamiliar setting or driven by learning.⁵

Finally, scholars have long been intrigued by chess as a game that is theoretically trivial but practically intractable. Chess is a finite, two-player, zero-sum game with perfect information (Neumann 1928).⁶ It is well known that, in any board configuration, either White has a winning strategy, Black possesses such a strategy, or both sides can guarantee themselves a draw (Zermelo 1913; Neumann and Morgenstern 1944). In the words of Osborne and Rubinstein (1994, p. 6):

The existence of such strategies suggests that chess is uninteresting because it has only one possible outcome. Nevertheless, chess remains a very popular and interesting game. [...] Even if White, for example, is shown one day to have a winning strategy, it may not be possible for a human being to implement that strategy. Thus while the abstract model of chess allows us to deduce a significant fact about the game, at the same time it omits the most important determinant of the outcome of an actual play of chess: the players' "abilities." Modeling asymmetries in abilities and in perceptions of a situation by different players is a fascinating challenge for future research, which models of "bounded rationality" have begun to tackle.

To this we only add that it is also important to *empirically* explore deviations from optimal play; and our study can be understood as taking a step in this direction.

2. Theory

Our analysis begins by developing a general model of satisficing with evaluation errors. After specializing the model to the case of chess and deriving its empirical predictions, we explain how to leverage object complexity to empirically distinguish between satisficing and standard forms of utility maximization.

⁵For conflicting evidence as to whether experience and skill in one strategic environment transfer to another one, see Palacios-Huerta and Volij (2008, 2009), Wooders (2010), and Levitt et al. (2010, 2011).

⁶There is a question as to whether chess should, in fact, be considered a finite game (see, e.g., the discussion in Osborne and Rubinstein 1994, p. 100). At least since 2014, the FIDE rule book specifies that a game is *automatically* drawn if the same position has occurred five times (cf. Article 9.6.1). Under official rules, chess is thus finite.

2.1. General Setup

A decision maker (DM) chooses a single object from every choice set $A \subseteq X$, where X is the grand set of all feasible alternatives. Each object $x \in X$ is characterized by a pair (v_x, σ_x) , where v_x denotes the true value of x to the DM and σ_x corresponds to its inherent complexity. The latter can be thought of as the length of an efficient representation of the object.

When the DM encounters an alternative x, she relies on her skill, S, to estimate v_x . By skill, we refer to the ability of the DM to break down complexity using her cognitive resources, experience, and so on. The outcome of the evaluation process is a noisy estimate:

(1)
$$u_x = v_x + \epsilon_x,$$

where the evaluation error ϵ_x is distributed normally with mean zero and standard deviation $\max\{0, \sigma_x - S\}$. Thus, when an object's inherent complexity is smaller than the DM's skill, the DM can correctly assess its true value. When object complexity exceeds the DM's skill, however, the DM needs to contend with noise in the evaluation process. The amount of noise increases in object complexity, and the DM's assessments are only correct in expectation.

In choosing from choice set A, the DM wishes to pick a high-value object while economizing on effort. To this end, she follows a satisficing procedure whereby she selects the first "satisfactory" alternative she encounters. Formally, the DM has in mind an aspiration level T. She randomly draws an object from A and evaluates it. The object is chosen if its estimated value exceeds T. Otherwise, the DM discards the object as non-satisfactory and draws the next one (without replacement). She continues in this fashion until she finds a satisfactory object or until she exhausts all alternatives in A without making a choice. In the latter case, the DM chooses the last object that she examined.⁷

The outcome of this choice procedure can be summarized by a random choice function C that maps every choice set A and every $x \in A$ to the probability C(x, A) of selecting x from the set A. Choice behavior is stochastic because object evaluations are noisy and because the evaluation order is random.⁸

⁷If the DM recalls the entire history of u's then she may choose the u-maximal object. This form of recall complicates the analysis to follow without adding substantive insights. Intuitively, this is because in the large sets we consider and assuming the threshold is not too high, exhausting all alternatives without exceeding T is a low probability event with a second-order effect on the theory's empirical predictions.

⁸The evaluation order does not favor any of the alternatives. Conditional on not stopping and not having been evaluated yet, every alternative x is equally likely to be examined next, and its conditional choice probability is determined by v_x and σ_x . If the evaluation order is allowed to favor some alternatives over others, then *any* random choice function C can be rationalized by a choice procedure that (i) draws the first alternative to be evaluated with its choice probability, and (ii) always chooses the alternative that is evaluated.

2.2. Application to Chess

In chess, the set X is the set of all possible legal moves of all pieces, and a choice set corresponds to all legal moves in some board configuration. Players' choice sets can thus be directly mapped to board configurations.

By Zermelo's Theorem, starting from any given configuration, either White can force a win, Black can force a win, or both sides can guarantee themselves a draw. It is therefore possible to associate every move in any board configuration with the ultimate outcome of the game under subsequent optimal play. A move that allows the DM to force a win yields positive payoff W, whereas a move that enables her opponent to do so produces a negative payoff L = -W. Moves that lead to draws generate a payoff of D = 0.9 Hence, $v_x \in \{W, D, L\}$ for every move x. As a matter of terminology, we say move x in some choice set is of type W, or simply a W-move, if $v_x = W$, with analogous designations for moves that yield payoffs of Dand L.

Our analysis focuses on choice sets that include at least one W- and at least one D- or L-move. In these sets, it is clear what it means to choose suboptimally. Moreover, the value of the maximal alternative remains unchanged as moves are added to the choice set. Both of these features are helpful in the empirical analysis.

We make two further assumptions. The first assumption guarantees that the DM stops and makes a choice whenever she encounters a W-move that she can correctly evaluate.

Assumption 1: The threshold T is between D and W.

The second assumption is only relevant for Predictions 1 and 4. It states that the inherent complexity of a D-move is not "too small" relative to that of L-moves. In Section 3, we argue that this assumption is plausible in the context of chess.

ASSUMPTION 2: In any choice set, the inherent complexity of any D-move is at least half of the inherent complexity of any L-move.

Under these assumptions, the model predicts that the choice probabilities of individual moves increase in their values.

PREDICTION 1: The choice probability of every W-move in a choice set exceeds the choice probability of every D-move in the same set, which in turn is larger than the choice probability of every L-move in the set.

⁹The assumed payoff structure reflects the fact that chess is usually considered a zero-sum symmetric game. This payoff structure is not critical for our analysis. For the predictions below, it suffices to assume that W > D > L and modify Assumption 2 appropriately.





Notes: Figure illustrates the evaluation process in our model, as applied to chess.

Figure 1 provides a graphical intuition for this and subsequent predictions. Area I in this figure corresponds to the probability that the DM discards a W-move, conditional on encountering one. Areas II and III correspond to the probability of accepting a D- or an L-move. Given that W > T > D and in light of the symmetric distribution of evaluation errors, areas I and II must each be smaller than one half. This observation, combined with random sampling of alternatives, means W-moves are chosen more frequently than D-moves. Figure 1 also makes clear why the inherent complexity of D-moves cannot be "too small" relative to moves of type L for Prediction 1 to hold. If it were, then area III may be larger than area II, which would imply that the choice probability of the corresponding L-move is larger than that of the corresponding D-move.

2.3. Object Complexity

The model makes two predictions regarding the connection between object complexity and choice behavior. The first prediction pertains to changes in the complexity of a single move, holding all else equal.

PREDICTION 2: As the complexity of a W-move increases, its choice probability decreases and that of all other available moves increases. As the complexity of an L- or D-move increases, its choice probability increases and all other choice probabilities decrease.

To build intuition, recall that higher object complexity corresponds to noisier evaluations of the respective moves and thus to "wider" distributions in Figure 1. Hence, an increase in the complexity of a *W*-move leads to more probability mass in area I, which implies a lower choice probability. As for the remaining moves' choice probabilities, they move in the opposite direction because the evaluation order is random, and in any given order, the choice probabilities of all subsequent moves must offset the change in the choice probability of the current move. For D- or L-moves, an increase in object complexity leads to more probability mass in area II or III, and thus to opposite comparative statics.

The second prediction compares choice probabilities within the same set.

PREDICTION 3: Fix a choice set A and two moves $x, y \in A$ with $\sigma_x < \sigma_y$. If both are W-moves then the choice probability of x is larger than that of y. However, if both are D- or L-moves then the choice probability of x is smaller than that of y.

Intuitively, a more complex W-move has more probability mass below the threshold than a simpler one, which implies that it is chosen with smaller probability. For two D- or two L-moves, larger complexity implies more mass above the threshold, and hence the opposite ranking of choice probabilities.

An immediate consequence of the discussion above is that mistakes become more likely as the complexity of any move in the choice set increases. By mistake, we mean that the DM picks a D- or an L-move in the presence of at least one W-alternative. Since we model the skill of the DM as equivalent to a reduction in the complexity of all moves, we also obtain the following corollary.

COROLLARY 1: Mistake probabilities decrease in the skill of the DM whenever her skill is lower than the largest object complexity in the choice set.

2.4. Satisficing vs. Maximization

Our satisficing-with-evaluation errors model combines two related but conceptually distinct ideas. First, we stipulate a noisy evaluation process that resembles the familiar randomutility framework in discrete-choice models (Luce 1959; Marschak 1960; McFadden 1974). An important difference between the evaluation process in our model and the standard discrete-choice setup is that ϵ_x in eq. (1) corresponds to an error in perception. Moreover, the variance of the error term depends directly on the object's inherent complexity and the DM's skill. This feature is essential for generating predictions on how object complexity affects choice.

The second idea is to incorporate noisy evaluations into a satisficing choice procedure. Drawing directly on Simon's (1955; 1972) discussion of bounded rationality and satisficing in chess, we model the DM as searching for and accepting the first "good enough" move. According to Simon (1955, p. 100), "in actual human decision-making, alternatives are often examined sequentially. [...] If a chess player finds a forced mate for his opponent, he generally adopts this alternative without worrying whether another alternative also leads to a forced mate."

Would the model yield similar predictions if it featured a standard maximization procedure rather than satisficing? By standard maximization, we mean any choice procedure in which

- the DM considers all moves in the choice set, obtains estimates of their values, and forms, for each move, a posterior belief about its value;
- the ranking of moves' expected values according to the posterior coincides with the ranking of the estimates; and
- the DM chooses the move with the highest posterior expected value.

Such a procedure encompasses a large class of models. Perhaps most importantly, it includes the standard " max_xu_x "–approach in the discrete-choice literature. It also includes some models of noisy cognition, in which agents choose the alternative with the highest Bayesian posterior mean (see, e.g., Woodford 2020 for a review).

Standard maximization and satisficing yield opposing predictions regarding the effect of object complexity on choice probabilities. Specifically, under maximization, the reverse of Prediction 3 holds.

PROPOSITION 1: Assume standard maximization as choice procedure, and consider any choice set A with two W-moves $x, y \in A$ such that $\sigma_x < \sigma_y$. The choice probability of x is smaller than that of y.

The intuition for Proposition 1 is as follows. Although the distribution of the difference $u_y - u_x$ is symmetric with mean zero, the distribution of u_y has "thicker" tails and, most importantly, a thicker right tail than that of u_x . The thicker right tail implies a positive probability of u_y being greater than u_x conditional on both of them exceeding "large" thresholds. Under maximization, the threshold they need to exceed is the maximum among all other moves and so, in expectation, it is large.

Proposition 1 and Prediction 3 provide a foundation for a new empirical test that distinguishes satisficing from utility maximization. This test is applicable in any environment in which there are at least two optimal alternatives that differ from one another in their inherent complexity. Satisficing predicts that the less complex optimal alternative will be chosen more often whereas maximization predicts the reverse.

3. Data

Taking the theory to the data requires information on actual choice behavior, as well as the available alternatives' values and object complexity. In this section, we describe how we



Figure 2: Screenshot of a Rated Game on Lichess

Notes: Figure shows a screenshot from a rated game between registered users on lichess.org. The green squares highlight the most recent move, i.e., $\triangle c4$.

assemble a data set on chess endgames that contains this information.

Information on choice behavior comes from lichess.org, one of the three most popular online chess platforms. 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 (see Figure 2 for a screenshot of a typical game). Although Lichess offers a choice between many different time limits, the majority of games that are actually hosted on the platform can be broadly classified as "speed chess," i.e., games in which players have significantly less time to complete their moves than under classical time controls.¹⁰ Lichess further distinguishes between casual and rated games. The latter determine player ratings and are therefore only available to registered users. In a nutshell, a player's strength rating increases (decreases) whenever she wins (loses) a rated game, and it increases (decreases) by more the stronger (weaker) her opponent was.¹¹ Anecdotal evidence from online messaging boards suggests that users often care intensely about their rating. 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.

We have data on the universe of rated games between human players from January 2013

¹⁰Some of our analyses restrict attention to games with classical controls, in which time pressure is less of an issue. As a robustness check, we also estimate alternative specifications that control for how much time a player has left on the clock. The results are qualitatively equivalent to those below (cf. Appendix C).

¹¹In calculating ratings, Lichess implements the Glicko-2 algorithm, which extends the better-known ELO system by incorporating a measure of uncertainty in approximately Bayesian fashion (Glickman 1999, 2000).

through August 2020. The available information includes players' usernames, ratings and realworld titles, if any (i.e., candidate master, national master, international master, grandmaster, etc.), the date and start time of the game, its outcome, and the precise sequence of moves.¹² We can therefore reconstruct all choice sets that a player faced as well as the moves she chose.

We augment these data with information on moves' types and, for W- and L-moves, their depth-to-mate (DTM). In principle, this information is computable for any legal move in any chess game, but doing so is computationally infeasible in the opening and middlegame phases. Our analysis thus focuses on endgame configurations with six or fewer pieces, which have been definitively solved by computer algorithms.¹³

These algorithms begin by constructing an exhaustive list of all possible (up to symmetry) legal board configurations with three chess pieces.¹⁴ Every configuration is examined, and the ones in which the player to move is in checkmate are stored as "mated in 0." Next, all configurations with the other side to move are evaluated. If one of them can reach a configuration that has previously been determined to be "mated in 0" by executing a legal move, then it is stored as "mate in 1." To find the set of configurations that are "mated in 2," the algorithm looks for configurations from which *all* possible legal moves lead to "mate in 1" configurations; and to determine configurations that are "mate in 3," it subsequently checks for configurations from which it is possible to directly reach a configuration that is known to be "mated in 2." Proceeding recursively, a configuration is classified as "mated in l" if and only if it is impossible to avoid moving to configurations that are "mate in w," with w = l - 1for at least one available legal move and $w \leq l-1$ for all others. By contrast, a configuration is marked as "mate in w" if it is possible to move to another one that is "mated in w - 1." This procedure continues until no further progress at classifying configurations is made, at which point all remaining configurations with three chess pieces are designated as "drawn." Essentially the same algorithm is next applied to board configurations with four pieces, then five, and then six.

The end result is a so-called tablebase in which *board configurations* are classified as either "drawn," "mated in l," or "mate in w." A particular *move* is said to be of type W with DTM d if it results in a new board configuration that, with the other player to move, is known to be "mated in d - 1." Thus, the minimal DTM among all available W-moves from any configuration that is "mate in w" is, by construction, equal to w. Similarly, a move is said to

¹²Lichess verifies that users do, indeed, hold the real-life titles they claim.

¹³Although board positions with seven pieces have, in principle, been solved as well, we do not rely on them in the analysis, because doing so would be impractical. Specifically, the open-source Syzygy tablebases do not contain information on DTM. The commercially available seven-piece Lomonosov tablebases do contain information on DTM but require about 140TB of storage. They are thus too large to be usable in most computing environments.

¹⁴It is not necessary to consider configurations with two lone kings, as these are automatically drawn.

Move Evaluation Move Evaluation Ŷ 8 ₩c6 W in 19 鬯b7 D7 W in 21 L in 34 ₩a8 Åh6 W in 47 L in 34 ₩a3 ₩c4 Ŵ Ī . 6 ₩a7 D ₩c8 L in 34 D ₩a5 ₿e2 L in 32 A 5 ₩a4 D₿e3 L in 32 4 ₩a2 D 當d3 L in 32 D鬯b6 L in 32 ₩a1 3 D ₿e1 L in 30 **鬯**b5 Ŷ D當d1 L in 28₩d3 2 ₩e2 DL in 28Åg6 1 ₩f1 D b f h с d е а g

Figure 3: Example of an Endgame Table

Notes: Figure provides an example of the information in endgame tablebases. The left panel shows the board configuration that is to be evaluated, assuming it is White's turn. Yellow-colored squares help visualize the set of available moves. The right panel shows the computer evaluation of each legal move, drawing on the Nalimov endgame tables. W, D, and L denote win-, draw-, and loss-moves from the perspective of the current player.

be an *L*-move with DTM *d* if it leads to a configuration that is "mate in d - 1." The maximal DTM among all *L*-moves from any configuration that is "mated in *l*" equals *l*. Moves that result in "drawn" configurations are classified as type D.¹⁵

For a concrete example of the content of such a tablebase, consider Figure 3. The left panel depicts the board configuration that is being examined, with the data for each available legal move shown on the right. The assessment of a move consists of two components: its type (i.e., W, D, or L), and, for a W- or L-move, its DTM. In this particular example, Wc6 corresponds to "W in 19," which means that, if White moves the queen to c6, then White can force checkmate in nineteen moves regardless of Black's response.

Endgame tablebases do not contain a measure of subgame depth for D-moves. To the best of our knowledge, there is no general approach to even identify D-moves by means other than elimination, which is how computer algorithms proceed. A classification-by-elimination approach, however, does not lend itself to a natural measure of complexity for D-moves other than suggesting it may be large relative to that of other alternatives in the choice set. While one could, in principle, calculate the depth of the subgame associated with a particular D-move by iteratively developing the game tree, doing so at scale is computationally infeasible. We therefore refrain from quantifying object complexity for D-alternatives, and restrict our empirical tests of Predictions 2 and 3 to W- and L-moves.¹⁶

¹⁵For additional information on algorithmic analysis of chess endgames, see, e.g., Thompson (1986).

¹⁶Prediction 1 relies on Assumption 2, which intuitively says that identifying a D-move with certainty is at least half as difficult as identifying L-moves. Per the discussion above, this assumption may be reasonable.

Given the results of extant computer analyses, the true value, or type, of essentially any endgame move is known with certainty, as is the DTM of W- and L-moves.¹⁷ We retrieve the relevant information for every legal move in every endgame configuration in the Lichess data by running several billion queries against the Syzygy and Nalimov tablebases (Nalimov et al. 2000; Man 2013).¹⁸

Our final sample contains nearly 227 million decision problems with a total of over 4.6 billion legal moves. There are four distinct sources of selection into this sample. First, we restrict attention to choice sets that contain at least one W- and at least one D- or L-move. We impose this restriction for expositional convenience. When a board configuration admits at least one W-move, any choice of a D- or L-alternative is unambiguously a mistake. In addition, we can distinguish "small" mistakes from "large" ones. The former correspond to picking a D-move in the presence of one or more of type W, while the latter consist of choosing an L-move instead. A disadvantage of this restriction is that the choice sets in our sample contain an above average share of W-moves.

The second and related source of selection pertains to how often different individuals reach a winning position, i.e., an endgame position with at least one W-move. The strongest players, for instance, may often mate their opponents before reaching the endgame stage. Similarly, very weak players may rarely be in a position to win endgames and might thus also be underrepresented in our sample. In what follows, we address this issue in two ways. First, whenever appropriate, we control for player fixed effects. Second, we reweight observations so that all players receive equal weight in our analysis. The results below should hence be interpreted as referring to a typical decision by the average player in our sample.¹⁹

Third, to minimize the risk that our findings are due to an unfamiliar setting or a lack of experience with similar decision problems, we exclude, for every player, the first one thousand endgame moves from winning positions. We are thus left with highly experienced DMs who are very familiar with the task at hand.

Finally, users on Lichess are not a random subset of all experienced chess players. In the

It is also reasonable to interpret our empirical test of Prediction 1 as a joint test of the prediction and Assumption 2.

¹⁷The only exceptions are positions with castling rights and configurations in which a lone king faces five other pieces. The former are extremely rare in endgames (< .01% of available legal moves in our data), while the latter are uninteresting (because 98.8% of available moves are of type W). Our empirical work excludes all board configurations for which information on DTM is not available.

¹⁸As a technical side note, the Syzygy tablebases do not contain information on DTM. In contrast to the Nalimov tables, they do, however, take into account the fifty-move stalemate rule. In rare instances, the fifty-move rule matters for correctly determining whether one player can unilaterally invoke a draw. We therefore rely on win-draw-loss information from the Syzygy database, while information on DTM comes from Nalimov's database.

¹⁹For robustness checks in which every individual observation receives the same weight, see Appendix C.

appendix, we address this potential source of concern by replicating our main results in an independent data set covering a large number of chess games in international tournaments. These data come from the online publication *The Week in Chess* (TWIC), which covers "all the latest news and games from international chess." The main disadvantage of this alternative data set is that there is less variation in the skill of players, and that it is more than two orders of magnitude smaller. These limitations notwithstanding, the TWIC data yield qualitatively similar conclusions (cf. Appendix D).

4. Descriptive Statistics

We next describe some of the basic features of our data, starting with the distributions of key variables. The upper two panels in Figure 4 display histograms for the total number of available legal moves in the DM's choice set (left) and for the number of W-, D-, and L-moves (right). Choice sets contain, on average, about 20.6 moves, of which 15.5 are W-moves. There is significant variation in choice-set size ranging from less than a handful of moves to more than thirty, or even fifty alternatives. Most, but by no means all, of this variation is due to differences in the availability of W-moves. The data therefore allow us to study choices from small, medium, and large sets.

The lower two panels of Figure 4 demonstrate that there is significant variation in our measure of object complexity. The lower-left panel depicts a histogram of the DTM of all available W- and L-moves. The modal DTM is 15 among the former and 24 among the latter. Given the size of the sample, there are millions of legal moves with DTM of five or less, and millions of moves with DTM of more than fifty. The lower-right panel plots the distribution of the minimal DTM among W-moves at the choice-set level. The histogram in this panel indicates that the data include choice sets in which evaluating some of the moves is straightforward, others where accurately classifying optimal moves likely exceeds the bounds of human cognition, and a great range of intermediate cases.

Figure 5 presents a binscatter plot of the relationship between moves' theoretical DTM and the total number of subsequent moves in the same game. The former may differ from the latter for a variety of reasons. For example, if the dominant player succeeds in mating her opponent but does not take the shortest path to victory, then the total number of subsequent moves may exceed the initial move's DTM. If, however, the losing player resigns or does not hold out as long as possible, then there will be fewer subsequent moves than implied by DTM. Importantly for our purposes, we observe a monotonically increasing relationship between moves' DTM and the length of the realized path of play. In other words, higher-DTM moves are, indeed, associated with longer subgames.

Some of our analysis focuses on suboptimal choice and mistakes. The summary statistics



Figure 4: Distributions of Key Variables

Notes: Panel (a) presents the distribution of the number of available moves, with panel (b) distinguishing between W-, D-, and L-moves. Panel (c) shows the distribution of DTM among W- and L-moves, and panel (d) depicts the distribution of the smallest DTM among all W-moves in the choice set.

Figure 5: Greater Depth to Mate is Associated with Longer Realized Game Paths



Notes: Figure shows binned scatter plots of the relationship between chosen moves' DTM (x-axis) and the total number of subsequent moves in the same game (y-axis). Hollow circles correspond to within-bin averages for all games, while solid circles restrict attention to games in which one of the players is eventually mated. The solid and dashed lines show the lines of best fit in the respective underlying move-level data.

in Table 1 show that about 6% of observed moves are mistakes. Given that, on average, W-moves account for about 75% of all available moves in a choice set, mistakes occur at nearly one quarter the rate one would expect if DMs were choosing at random. There is thus evidence of significant skill among the players in our data. Yet, the raw data also imply that mistakes occur with a certain regularity. They are not rare events. Moreover, mistakes tend to be consequential. A player whose current move is a mistake is about 43 p.p.—or roughly 58%—less likely to ultimately win the game than one who chooses a W-move. At the same time, the probability of a loss more than doubles.²⁰

To study choices at a more granular level, we partition the data into classes of choice sets so that two sets are in the same class if and only if they have the same number of W-, D-, and L-moves. We observe 9,701 such classes with more than a hundred choices. For each class, we calculate the empirical frequency with which a legal move of a particular type is chosen, and plot the observed differences in Figure 6.

The top-left panel shows a histogram of differences in choice frequency between D- and W-moves—one estimate for each class—and the bottom-left panel does so for the difference between D- and L-alternatives. To the right of each histogram, we show the CDF of the associated p-values. Under the null hypothesis that choice probabilities do not depend on moves' types, the p-values of the estimated differences should be approximately uniformly distributed over the unit interval. This is not the case. In the vast majority of classes of choice sets, different types of moves are chosen with frequencies that are statistically distinguishable at conventional significance levels. More importantly, the signs of the observed differences are generally as postulated by the theory. That is, W-moves are more likely to be chosen than D-moves in over 99% of classes, and D-moves are more likely to be chosen than L-moves in over 93% of classes. Prediction 1 is thus borne out in the data.

²⁰Mistakes do not always result in forgone wins because the player's opponent may subsequently also choose a suboptimal move. Similarly, due to the potential for mistakes in future moves, choosing a W-move in a particular instance does not automatically mean that the player will ultimately win the game.

				Perc	entile		
Variable	Mean	SD	25%	50%	75%	95%	N
A. Move Characteristics							
<i>Type:</i> W. Morro	0.60	0.46					4 617 441 579
D Move	0.09	0.40					4,017,441,575
L Move	0.23	0.42					4,017,441,575
DTM.	0.08	0.27					4,017,441,575
W Morros	25.80	17 99	19	1 2	9 9	50	2 157 979 209
<i>u</i> -Moves	20.89 30.35	13.46	13 22	23 28	36	59 50	296 522 573
	50.55	10.40	22	20	50	00	230,022,010
B. Choice-Set Characteristics							
Set Composition:							
Total Number of Legal Moves	20.58	10.35	13	20	28	38	$226,\!955,\!095$
Number of W -Moves	15.48	11.09	6	15	24	34	$226,\!955,\!095$
Number of <i>D</i> -Moves	3.81	4.30	1	2	5	13	$226,\!955,\!095$
Number of <i>L</i> -Moves	1.29	2.73	0	0	2	7	$226,\!955,\!095$
DTM of W-Moves:							
Shortest	20.91	16.18	9	17	29	51	$226,\!955,\!095$
Median	25.81	17.34	13	22	33	59	$226,\!955,\!095$
Longest	31.62	20.05	17	27	39	69	$226,\!955,\!095$
DTM of L-Moves:							
Shortest	26.95	11.39	20	26	34	44	77,045,730
Median	30.19	12.50	22	28	36	48	77,045,730
Longest	34.27	16.31	24	32	40	60	77,045,730
C. Outcomes							
Mistakes:							
Any Type of Error	0.06	0.24					226,955,095
Small Mistake	0.05	0.23					226,955,095
Large Mistake	0.01	0.08					226,955,095
Result of Game:							, ,
If Current Move is Mistake:							
Win Game	0.31	0.46					13,052,773
Draw	0.49	0.50					13,052,773
Lose Game	0.19	0.40					13,052,773
If Choose W-Move:							
Win Game	0.74	0.44					213,902,322
Draw	0.20	0.40					213,902,322
Lose Game	0.05	0.22					213,902,322
D. Timing							
$\frac{D. \text{ Time Left on Clock (in sec.)}}{\text{Time Left on Clock (in sec.)}}$	72.35	192.87	8	22	69	296	212 295 223
Deliberation Time (in sec.)	1.66	3.07	0	1	2	5	212,249,738
E. Player Characteristics	0.50	a			a c==	0	005
Total Number of Endgame Moves	2,584	2,469	1,297	1,793	2,877	6,696	237,232
Average Rating	1,733	281	1,533	1,716	1,917	2,222	237,232
Real-World Title	0.01	0.10					237,232

 Table 1: Summary Statistics

Notes: Table displays summary statistics for selected variables in the data. Each observation in panel A corresponds to a legal move, and observations in panels B–D correspond to decision problems. Panel E contains player-level information. Observations are reweighted so that all players receive equal weight. The number of observations related to the timing of moves is smaller because the raw data do not include this information for games that were played prior to April 2017.



Figure 6: Comparing Choice Frequencies of W-, D-, and L-Moves

mean differences in choice frequencies of D- and W-moves in choice sets with identical composition of moves, conditional on observing at least 100 choices from sets with that makeup. Negative estimates imply that, in the respective classes of choice sets, D-moves are, on average, less frequently chosen than W-alternatives. Panel (b) depicts the empirical CDF of the two-sided p-values associated with the estimated differences in panel (a). It also shows results from a Kolmogorov-Smirnov test against the null hypothesis of a uniform distribution of p-values. Panel (c) mirrors panel (a) but focuses on the difference in choice frequencies between D- and L-moves. Positive estimates imply that D-moves are, on average, more frequently chosen than L-alternatives. Panel (d) presents the empirical CDF of p-values for the null hypothesis that, in choice sets with a particular composition, L-moves are as frequently chosen as D-moves. All p-values account for two-way clustering by player and game.

5. Object Complexity and Choice Frequencies

Predictions 2 and 3 relate object complexity to choice probabilities. Prediction 2 is concerned with how choice probabilities change as, all else equal, evaluating a particular legal move becomes more difficult. Because every board configuration is associated with a fixed choice set, the relevant empirical comparisons are necessarily across decision problems. By contrast, Prediction 3 compares the choice probabilities of different moves of the same type within the same choice set. According to Prediction 3, we should see that the W-move with the smallest DTM is chosen more frequently than the W-alternative with the second-smallest DTM value, which, in turn, is picked more often than that with the third-smallest value, and so on. For L-moves, the opposite pattern should emerge.

5.1. Across-Choice-Set Comparisons

To investigate Prediction 2, we focus on W- and L-moves and estimate the following econometric model:

(2)
$$Choose_a = \omega DTM_a \times W_a + \iota DTM_a \times L_a + \pi W_a + \phi_{A\backslash a} + \mu_p + \varepsilon_a.$$

Here, $Choose_a$ is an indicator for whether player p facing choice set A picked move a, DTM_a denotes that move's depth to mate, W_a and L_a are indicators for its type, and μ_p is a player fixed effect. We also include $\phi_{A\setminus a}$, a fixed effect for the other moves in the choice set. In constructing this fixed effect, we assume that, in line with the theory, moves can be reduced to their types and inherent complexity. Since we do not measure object complexity for D-moves, $\phi_{A\setminus a}$ conditions (only) on the vector of DTM values for W- and L-moves and the type composition of the choice set, i.e., the number of W-, D-, and L-alternatives. By including $\phi_{A\setminus a}$, we aim to approximate the thought experiment in which we vary the object complexity of one particular alternative, holding the values and object complexity of all other moves fixed.

Table 2 shows results from estimating variants of the regression model in eq. (2). In the first three columns, we study W- and L-moves from all board configurations. In the last three columns, we restrict attention to choice sets that do not contain any D-moves. The idea that our fixed effects appropriately control for the object complexity of all other moves is most plausible in the latter set of specifications. Reassuringly, the results in both sets of columns are qualitatively similar. Regardless of which sample we consider, we find that individual W-moves are significantly less likely to be chosen as object complexity increases. By contrast, the choice probabilities of L-moves increase in their inherent complexity.

The estimates in Table 2 are not only directionally consistent with Prediction 2, but are

	(1)	(2)	(3)	(4)	(5)	(6)
			Probability of C	Choosing Move	9	
	W-Moves	<i>L</i> -Moves	W- & L -Moves	$W ext{-Moves}$	L-Moves	W- & L-Moves
DTM \times W-Move (\div 100)	-1.140 (0.002)		-0.974 (0.002)	-0.521 (0.004)		-0.333 (0.003)
DTM \times <i>L</i> -Move (÷ 100)	()	0.034 (0.001)	0.007 (0.001)	()	0.066 (0.002)	0.054 (0.002)
W-Move (\div 100)		()	63.450 (0.069)			41.857 (0.114)
Fixed Effects: Number of W - \times D - \times L -Moves						
\times Complexity of Other Moves Player	Yes Yes	Yes Yes	Yes Yes	Yes Yes	Yes Yes	Yes Yes
Board Configurations	All	All	All	$\mid D \mid = 0$	$\mid D \mid = 0$	$\mid D \mid = 0$
Mean of LHS Variable (%)	9.397	0.785	8.495	10.881	1.183	8.078
R^2 N	0.372 3,457,878,398	$0.209 \\296,522,573$	0.368 3,754,400,971	0.513 398,856,135	0.209 111,905,262	0.494 510,761,397

Table 2: Choices as a Function of DTM

Notes: Entries are coefficients and standard errors from estimating ω , ι , and π in variants of eq. (2) by ordinary least squares. All regressions control for player fixed effects as well as a fixed effect for the combination of the number of W-, D-, and L-moves and the vector of DTMs of all other W- and L-moves in the same choice set, as explained in the text. The unit of observation in each regression is an available legal W- or L-move, and observations are reweighted so that all decision problems for a particular player and all players receive equal total weight in the full sample of moves. The sample in cols. (1)–(3) includes all board configurations in our data, whereas cols. (4)–(6) restrict attention to configurations for which the associated choice sets do not contain D-moves. 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 two-way clustered by player and game, and are shown in parentheses.

also economically large. According to the coefficients in cols. (1) and (4), increasing DTM by two nodes decreases the choice probability of the corresponding W-move by 1.0–2.3 p.p. This decrease translates to a 10%–24% reduction relative to the mean choice probability of W-moves. The evidence in this table therefore suggests that object complexity is an empirically meaningful determinant of choice.

5.2. Within-Choice-Set Comparisons

Our examination of Prediction 3 begins by looking at the raw data. We first partition choice sets into classes according to the number of available W-moves. In the left panel of Figure 7, we then plot, for each class, the empirical frequency with which DMs choose the W-move with the smallest DTM, the second-smallest DTM, and so on. We observe large and precisely estimated differences in choice frequencies. In line with Prediction 3, these frequencies decrease monotonically in moves' relative complexity, and all differences are statistically significant at conventional levels.

The right panel of Figure 7 repeats this analysis for L-moves. Choice frequencies for L-moves are much smaller and less precisely estimated. This is partly due to the relatively low number



Figure 7: Choice Frequencies, by Relative Object Complexity

Notes: Figure compares choice frequencies of W- and L-moves according to the ordering of their DTMs relative to other moves of the same type in the same choice set. Panel (a) does so for W-moves, and panel (b) focuses on L-moves. Each plot within either panel restricts attention to choice sets that contain a particular number of W- or L-moves with distinct DTM values. Error bars correspond to 95%-confidence intervals and account for two-way clustering by player and game.

of choice sets with more than three L- alternatives (cf. Figure 4). Comparing mean choice frequencies within the same class, we find that out of fifty-six possible pairwise comparisons, over 80% have the correct sign according to Prediction 3 and over 55% are statistically significant at the 5%-level. Only three comparisons, corresponding to about 5% of the total, have the wrong sign and are significant at the 5%-level. Given that some false positives are expected even if the null hypothesis holds, we conclude that the data are broadly consistent with the theory.

Under the maintained assumption of noisy evaluations, the ranking of choice frequencies for W-moves in the left panel of Figure 7 is inconsistent with standard maximization (cf. Proposition 1). This is perhaps not surprising. First, the majority of games in our data are best thought of as "speed chess," and players may simply not have enough time to seriously consider all available moves. Second, some of the choice sets in the data are so large that DMs might find it challenging to evaluate all, or even most, alternatives. Third, the data include choice sets from which skilled-enough players may instantly recognize and gravitate towards very simple moves. For all of these reasons, we view the above findings as a "proof of concept" for our empirical test. It indeed distinguishes satisficing from maximization in an environment where satisficing is most likely to occur. It is entirely possible that choice behavior exhibits different patterns when choice sets are small, time pressure is not an issue, or none of the available moves are simple.

To investigate this possibility, we compute within-choice-set correlations between choice frequencies and moves' relative complexity in such settings. Table 3 presents our estimates, which are based on the following econometric model:

(3)
$$Choose_a = \beta Percentile Rank_a + \chi_{A,p} + \epsilon_a,$$

where $Choose_a$ denotes an indicator for whether a player p facing choice set A picked move a, $\chi_{A,p}$ is a fixed effect for the particular decision problem of that player (i.e., the task of choosing a move from A), and *Percentile Rank_a* corresponds to the respective move's DTM-based percentile ranking among all available moves of the same type. By construction, this variable is equal to zero for the simplest W- and L-moves in each choice set, and one for the most complex ones.²¹

The coefficient of interest is β . Under the parametric assumptions behind eq. (3), β corresponds to the average difference in choice probability between the simplest and most complex moves of the same type. Since our empirical specification controls for $\chi_{A,p}$, all identifying variation comes from comparing alternatives within the same choice set.

Cols. (1A) and (1B) of Table 3 confirm that, in the full sample, choice frequencies and relative complexity are negatively correlated among W-moves, and positively correlated among L-moves. Taking the point estimates at face value, the simplest winning move in a choice set is about 30 percentage points more likely to be chosen than its most complex counterpart, while the simplest losing move is about 0.3 percentage points less likely to be selected than the most complex one. Relative to the respective mean choice frequencies, the coefficients in cols. (1A) and (1B) are very large.

The next three columns restrict attention to settings that meet one of the following criteria: (i) long time controls (so that each player has, in expectation, at least twenty-five minutes for deliberation per game); (ii) small choice sets (with ten or fewer moves); (iii) none of the available alternatives are easily recognizable as good (because minimal DTM exceeds fifty). Consistent with satisficing, we continue to find that W-moves' choice frequencies are strongly negatively correlated with their complexity. For L-moves, the opposite pattern holds.

The last column presents estimates of β , excluding the simplest move of a given type in every choice set. Although the coefficient in col. (5A) does decrease in magnitude, it remains very large relative to the mean choice frequency of the remaining moves. Thus, Prediction 3 continues to hold.

Going beyond aggregate correlations, Figure 8 examines the relationship between choice

²¹Formally, Percentile Rank_a $\equiv \frac{r_a-1}{M_A-1}$ where r_a corresponds to move *a*'s ordinal DTM-based rank among moves of the same type in choice set *A* (simplest = 1, second-simplest = 2, and so on), and M_A denotes the total number of moves of that type. Using moves' percentile rank enables us to sidestep issues related to the fact that different choice sets contain different numbers of *W*- and *L*-moves. Appendix Table AT.1 presents results based on ordinal rather than percentile rank. Both sets of results are qualitatively equivalent.

Panel A: W-Me	oves				
_		Prob	ability of Choosing	Move	
	(1A)	(2A)	(3A)	(4A)	(5A)
Percentile Rank (\div 100)	-30.147 (0.017)	$-32.921 \\ (0.116)$	-54.851 (0.044)	-23.424 (0.064)	-9.009 (0.007)
Fixed Effects: Player \times Choice Set	Yes	Yes	Yes	Yes	Yes
Sample	Full	Long Time Controls	≤ 10 Total Moves	$\mathrm{DTM} \geq 50$	Excl. Simplest Move
Mean of LHS Variable $(\%)$	10.305	11.129	26.210	14.992	4.301
$\frac{R^2}{N}$	0.234 3,435,257,516	$0.246 \\92,878,586$	$0.319 \\ 148,599,747$	$0.208 \\ 89,336,853$	$0.133 \\ 2,986,665,574$

Table 3:	Within-Choice-Set	Correlations	between	Relative	Complex	xity and	Choice	Frequ	uencies
1 0010 01	filling offered bee	0.011.0100010110	~~~~~	1001000100	00110101	and, come	0110100		10110100

		Prob	ability of Choosing	Move	
	(1B)	(2B)	(3B)	(4B)	(5B)
Percentile Rank (\div 100)	$0.266 \\ (0.005)$	$0.148 \\ (0.019)$	$1.308 \\ (0.030)$	$0.189 \\ (0.051)$	$0.576 \\ (0.012)$
Fixed Effects: Player \times Choice Set	Yes	Yes	Yes	Yes	Yes
Sample	Full	Long Time Controls	≤ 10 Total Moves	$DTM \ge 50$	Excl. Simplest Move
Mean of LHS Variable $(\%)$	0.452	0.188	1.352	0.588	0.498
$\frac{R^2}{N}$	0.277 277,507,532	$0.276 \\ 8,300,536$	$0.306 \\ 34,762,492$	$0.344 \\ 4,095,758$	0.270 175,675,645

Panel B: L-Moves

Notes: Entries are coefficients and standard errors from estimating β in eq. (3) by ordinary least squares. The upper panel restricts attention to different samples of W-moves, while the lower panel focuses on L-alternatives. As explained in the text, all regressions control for a decision problem fixed effect, i.e., a fixed effect for a particular player and the set from which she chooses. The unit of observation in each regression is an available legal W- or L-move, and observations are reweighted so that all moves for a given player and all players receive equal total weight. All estimates are scaled to correspond to the percentage-point change in choice probability associated with a one-unit increase in the respective move's percentile rank, as defined in the text. Standard errors are two-way clustered by player and game, and are shown in parentheses.

frequencies and relative complexity at the level of individual players. To give maximization its best shot, we restrict attention to decision problems that simultaneously meet criteria (i), (ii), and (iii) above. In these settings, players have time to consider all available alternatives, none of which are easily recognizable as good. Focusing on W-moves in this highly selected sample, we estimate the regression model in eq. (3) for every player for whom we observe at least ten such decisions.

The left panel of Figure 8 plots the distribution of the resulting point estimates. For over 90% of players, choice frequencies are negatively correlated with W-moves' relative complexity. The right panel of Figure 8 shows the empirical CDF of the one-sided *p*-values associated with the estimates on the left. The relevant *p*-values are one-sided because the null hypothesis of noisy maximization implies that $\beta \geq 0$ (cf. Proposition 1). Under this null, the



Figure 8: Testing Noisy Maximization at the Player Level

Notes: Figure presents player-level tests of Proposition 1. Panel (a) plots histograms of estimates of β in eq. (3), restricting attention to W-moves in decision problems in which a player chooses from ten or fewer alternatives with minimal DTM among W-moves of at least fifty, subject to classical or less restrictive time controls. We estimate $\hat{\beta}$ separately for all players for whom we observe at least ten such decision problems. Estimates are scaled to be directly comparable to their counterparts in the upper panel of Table 3. Panel (b) shows the empirical CDF of the one-sided *p*-values associated with the point estimates in panel (a). It also shows results from a Kolmogorov-Smirnov test against the null hypothesis of a uniform distribution of *p*-values. All *p*-values account for clustering across moves in the same choice set.

distribution of p-values should first-order stochastically dominate the uniform distribution.²² This, however, is not what we observe. The actual distribution of p-values is first-order stochastically dominated by the uniform distribution and a formal Kolmogorov-Smirnov test rejects the limit case of uniformity at the 99%-confidence level.

We conclude that simpler winning moves tend to be chosen more frequently than more complex ones, even in settings in which players are *a priori* most likely to exhibit maximizationbased behavior. We can thus reject noisy maximization as the primary choice procedure in the data.

6. Moderators of Object Complexity

In the model, skill helps DMs offset the adverse effect of object complexity. An increase in skill is analogous to a decrease in complexity for all moves in the choice set. Therefore, skilled players are less likely to make sub-optimal choices than unskilled ones (cf. Corollary 1).

We test this prediction by distinguishing between titled and untitled players. About 1.5% of decisions in our data are made by users who hold real-world chess titles, such as candidate master, national master, international master, or grandmaster. These are some of the best

²²That is, we should have $Pr(p \le \alpha \mid H_0) \le \alpha$ for all $\alpha \in [0, 1]$. This claim follows from Definition 8.3.26 and Theorem 8.3.27 in Casella and Berger (2001). The intuition behind it is as follows. If $\beta = 0$, then the observed *p*-values should be uniformly distributed over the unit interval. If $\beta > 0$, however, then we would expect to see fewer small (one-sided) *p*-values and more large ones, implying first-order stochastic dominance.

	(1)	(2)	(3)	(4)
		Probability	of Mistake	
Titled Player (\div 100)	-0.813 (0.048)		$-0.179 \\ (0.061)$	
Other Title $(\div 100)$		-0.768		-0.155
Grandmaster (\div 100)		$egin{array}{c} (0.051) \ -1.134 \ (0.130) \end{array}$		$egin{array}{c} (0.064) \ -0.350 \ (0.197) \end{array}$
Hypothesis Tests (<i>p</i> -value): H_0 : No Differences between Players H_0 : Grandmasters = Other Titled Players	< 0.001	$< 0.001 \\ 0.009$	0.004	$\begin{array}{c} 0.012\\ 0.344\end{array}$
Fixed Effects: Number of W - \times D - \times L -Moves \times Complexity of Moves	Yes	Yes	Yes	Yes
Mean of LHS Variable (%)	6.039	6.039	2.748	2.748
Board Configurations	All	All	$\mid D \mid = 0$	$\mid D \mid = 0$
$\frac{R^2}{N}$	0.366 226,955,095	0.366 226,955,095	$0.410 \\ 27,600,514$	$0.410 \\ 27,600,514$

Table 4: Titled vs. Untitled Players' Frequency of Mistakes

Notes: Entries are coefficients and standard errors from estimating τ in variants of eq. (4) by ordinary least squares. As explained in the text, all regressions include a fixed effect for the combination of the number of W-, D-, and L-moves and the vector of all DTM values of the W- and L-moves in the choice set. The unit of observation in each regression is a decision problem, and individual observations are reweighted so that all players receive equal total weight. The sample in cols. (1) and (2) includes all board configurations in our data, whereas cols. (3) and (4) restrict attention to configurations for which the associated choice sets do not contain D-moves. All estimates are scaled to correspond to the percentage-point change in the probability of an error associated with a one-unit increase in the respective regressor. Standard errors are two-way clustered by player and game, and are shown in parentheses.

chess players in the world, many of whom are professionals. By contrast, untitled users in our sample are best thought of as experienced hobbyists. It is thus plausible that titled players are, on average, more skilled than untitled ones. Given the sheer size of the data, we observe enough decisions by either type of player to ask whether their mistake frequencies are systematically different, and, if so, how skill and object complexity interact.

We study the first of these questions by estimating the specification

(4) $Error_d = \tau Titled_p + \xi_A + \varepsilon_d,$

where $Error_d$ indicates whether decision d from choice set A by player p is a mistake, and $Titled_p$ indicates whether player p holds a title of candidate master or higher. We control for the difficulty of players' decision problems by adding a fixed effect, ξ_A , for the number of L-, D-, and W-moves as well as the vector of DTM values for all L- and W-moves.

Relying on the regression specification above, the results in Table 4 imply that titled players are, on average, significantly less likely to make mistakes than untitled ones. Cols. (1) and (2) show fewer mistakes for all board configurations, while cols. (3) and (4) focus on choice sets that do not contain D-moves. The estimates in cols. (2) and (4) even suggest that



Figure 9: Titled vs. Untitled Players' Frequency of Mistakes, by Minimal DTM among W-Moves

Notes: Figure compares mistake frequencies for titled and untitled players by minimal DTM among available W-moves. Panel (a) plots raw mistake frequencies, while panel (b) shows point estimates and 95%-confidence intervals for ϕ_i in eq. (5), i.e., the difference in error rates between titled and untitled players when the length of the shortest winning path equals *i*. Negative point estimates indicate fewer mistakes among titled players. Confidence intervals account for two-way clustering by player and game.

grandmasters are less likely to err than other titled players facing similar choice sets. In chess, grandmasters are typically regarded as the "best of the best," but they only account for about 0.15% of decisions in our data, which may explain why the difference between them and other titled players is only statistically significant in col. (2). The latter fact notwithstanding, Corollary 1 is borne out in the data.

To address the question of how skill and object complexity interact, Figure 9 studies whether mistake frequencies for titled and untitled players depend on the minimal DTM among Wmoves in the choice set. From a theoretical point of view, minimal DTM corresponds to the minimal skill level that is required to precisely identify the value of at least one W-move in the choice set. Empirically, the correlation between minimal DTM and other low-dimensional proxies for object complexity, such as the average or median DTM, exceeds 0.9.

The left panel of Figure 9 plots raw mistake frequencies for titled and untitled players. For minimal DTM values below thirty, we observe very small, if any, differences between both sets of players. Only for choice sets in which evaluating W-moves is arguably difficult do we see that titled players commit significantly fewer mistakes.

A potential issue with the raw data is that titled and untitled players might be choosing from sets that are systematically different, even conditional on the same minimal DTM. The right panel of Figure 9 addresses this concern by estimating the difference in the frequency of mistakes controlling for the number of W-, D-, and L-moves as well as the DTM values of all W- and L-moves. Specifically, the results in this panel correspond to $\{\phi_i\}$ in the following





Notes: Figure compares mistake frequencies across different sets of time controls by minimal DTM among available W-moves. Panel (a) plots raw mistake frequencies, while panel (b) shows estimated mean differences relative to long time controls after conditioning on a fixed effect for the number of available moves of each type as well as the vector of DTM values for all L- and W-moves. Confidence intervals account for two-way clustering by player and game.

regression:

(5)
$$Error_d = \sum_{i=1}^{120} \phi_i \lfloor DTM \rfloor_{i,A} \times Titled_p + \xi_A + \varepsilon_d,$$

where $Error_d$, $Titled_p$, and ξ_A are defined as in specification (4) above; and $\lfloor DTM \rfloor_{i,A}$ is an indicator equal to one if the minimal DTM among all W-moves in choice set A is equal to i.

The regression estimates accord with the model predictions. At the risk of overinterpreting the patterns in the data, it appears that cognitive constraints begin to negatively affect untitled players' decision-making when minimal DTM is between twenty and thirty.

Skill is the only factor that moderates complexity in the model. In reality, there might be additional moderators. For instance, DMs may examine alternatives in less depth when afforded less time to make choices. If so, we would expect the frequency of mistakes to vary with games' time controls.

On Lichess, games are played subject to various time controls, which govern how fast decisions need to be made. For the sake of parsimony, we broadly categorize the available options as follows: (i) ultra-short time controls (UltraBullet and Bullet), which result in matches that conclude within a few minutes; (ii) short time controls (Blitz and Rapid), under which games have an expected duration of up to fifty minutes; and (iii) long time controls (Classical and Correspondence), resulting in matches that can last hours, or even days.

As expected, Figure 10 shows that players are less likely to make mistakes when they



Figure 11: Mistake Frequencies, by Skill Level and Time Controls

Notes: Figure compares the frequency of mistakes for titled and untitled players under different time controls. Panel (a) plots raw frequencies, while panel (b) shows estimated mean differences between both sets of players, controlling for the number of available moves of each type as well as the vector of DTM values for all *L*- and *W*-moves. Negative values correspond to fewer mistakes among titled players. Moving from the left to the right within each panel, time constraints on players' decision-making become less and less restrictive. Error bars correspond to 95%-confidence intervals and account for two-way clustering by player and game. Braces indicate which estimated differences are being tested for equality.

are given more time to make a decision. Importantly, the gap in the frequency of mistakes across different time controls tends to become wider as the minimal DTM among W-moves increases. This observation is borne out in the raw data (left panel) as well as in regression estimates analogous to those in Figure 9 (right panel). In sum, the evidence in Figures 9 and 10 suggests that both skill and time moderate the adverse impact of complexity on the quality of decision-making.

This finding suggests that, in our model, the standard deviation of the evaluation error ϵ_x might be more appropriately specified as $\max\{0, \sigma_x - f(S, T)\}$, where T denotes time, and $\frac{\partial f}{\partial S}, \frac{\partial f}{\partial T} > 0$. Going one step further, we can ask whether $\frac{\partial^2 f}{\partial S \partial T} > 0$. That is, are time and skill complements?

The evidence in Figure 11 suggests that the answer is "yes"—at least in our setting. To arrive at this conclusion, we implement a simple difference-in-differences approach. That is, we compare the difference in the frequency of mistakes between titled and untitled players across different sets of time controls. The left panel of Figure 11 shows that, in the raw data, the gap between titled and untitled players widens as both are afforded more time to make decisions. The right panel presents regression adjusted estimates for the difference between both sets of players controlling, as above, for the number and types of all available moves as well as the vector of DTM values. Reassuringly, the numbers in either panel suggest that titled players benefit more from longer time controls than untitled ones. Thus, time and skill appear to be complements.

7. Complexity at the Set Level

Moving beyond object complexity and its moderators, there are likely other, set-level factors that complicate choice. For instance, a large literature in psychology argues that choice set size adversely affects decision-making (see, e.g., Iyengar and Lepper 2000; Schwartz 2004; Chernev et al. 2015). In the context of chess, Anderson et al. (2017) show that players are more likely to blunder when choosing from sets that contain a smaller share of W-moves.²³ In this section, we build on their work. We first return to the model and develop its predictions on how the size and composition of the choice set affect mistake probabilities. We then verify these predictions in the data. Finally, we propose a definition of *composition complexity* that may be useful in other applications.

7.1. Model Predictions

In the model, every object is characterized by its type and inherent complexity. Our analysis thus far studied object complexity. We now focus on moves' types.

Consider adding a move to the choice set. Random sampling together with noisy evaluations imply that such an addition leads to choice probability shifting away from all previously available moves toward the new one. Thus, if the added alternative is an L- or D-move, the probability of making a mistake increases. If, however, the added move is a W-move then the probability of making a mistake decreases. Comparing the effect of adding a D- to that of adding an L-move, the model predicts that mistake probabilities increase by more when a D-move is added to the choice set.

PREDICTION 4: For any choice set, adding a D-move to the set increases the probability of a mistake by more than adding an L-move.

Returning to Figure 1, the intuition for Prediction 4 is that, given any evaluation order, moving from area III to II results in a greater probability of accepting a suboptimal move.

Consider next replacing a W-move with an L- or D-alternative. Because this replacement can be accomplished by first adding an L- or D-move to the set and then removing the W-move, mistake probabilities increase. A more subtle change is to replace an L-move with a D-move. By Prediction 4, this change in the composition of suboptimal alternatives also increases mistake probabilities. Broadly summarizing, our model predicts that what complicates decision-making is not necessarily the size of the choice set but its composition.



Figure 12: Mistakes as Function of the Number and Type of Available Moves

Notes: Panel (a) compares observed mistake rates for choice sets of a particular size with the frequency of mistakes that would be expected if players chose a move at random. Panel (b) shows mistake rates for different shares of W-, D-, and L-moves.

7.2. Empirical Evidence

Drawing on the raw data, the left panel of Figure 12 depicts the relationship between the size of the choice set and the frequency of mistakes. It also contrasts observed mistake rates with the probability of making a mistake when choosing a move at random. There is a clear negative association between the number of available alternatives and the frequency of mistakes. However, because larger choice sets turn out to contain, on average, a greater share of W-moves, the observed decline in the frequency of mistakes might be mechanical. Measured as a fraction of the baseline risk of making a mistake when choosing at random, error rates remain nearly constant over approximately 85% of our sample. Only for choice sets with more than thirty moves do we observe some increase in relative risk. Taken at face value, the evidence in this panel is difficult to reconcile with elementary notions of choice overload, according to which larger choice sets are *per se* more complex.

What appears to affect mistake frequencies is the mix of moves. The right panel of Figure 12 illustrates this point by examining how the rate of mistakes varies with the relative frequency of each type of move in the choice set. The entry on the bottom left, for instance, indicates that when the fraction of W- and D-moves ranges from zero to ten percent each, about 16% of choices end up being mistakes. Comparing rows within the same column, we see mistakes become more frequent as W-moves become relatively more scarce. And reading across columns of the same row, mistakes also tend to increase as D-moves become more frequent. In the bottom row, for instance, the frequency of mistakes increases from 16% to

 $^{^{23}}$ Anderson et al. (2017) do not distinguish between D- and L-alternatives, and they do not consider moves' DTM in any of their analyses.

more than 30%. Since comparisons within the same row hold the fraction of *W*-moves—and hence the mechanical probability of making a mistake—approximately fixed, the data suggest the quality of decision-making correlates with the *composition* of the remaining, inferior alternatives.

To more rigorously establish that different types of changes to the choice set affect mistake frequencies as predicted by the theory, we estimate η in the following econometric model:

(6)
$$Error_{d} = \eta Number Moves_{A}^{t} + \kappa_{A}^{\lfloor DTM \rfloor} + \theta_{A}^{-t} + \mu_{p} + \varepsilon_{d}$$

In this specification, $Error_d$ is an indicator for whether decision d of player p choosing from choice set A is a mistake, and $Number Moves_A^t$ denotes the number of available legal moves of type $t \in \{W, D, L\}$. Although not required by the theory, we account (in a low-dimensional way) for the contribution of object complexity to mistake frequencies by including $\kappa_A^{\lfloor DTM \rfloor}$, a fixed effect for the minimal DTM among all available W-moves.²⁴ The fixed effect θ_A^{-t} controls in different ways for the number of moves in the same choice set that are not of type t. As we vary θ_A^{-t} across columns in Table 5, the interpretation of η changes to closely align with the different changes to choice-set composition that we discussed above. For example, col. (6) conditions on the choice-set size and the number of W-moves, so that $\hat{\eta}$ should be interpreted as the change in mistake probabilities when an L-move is replaced with one of type D. By including μ_p , the model in eq. (6) also controls for player fixed effects.

The regression in col. (1) of Table 5 conditions on the number of D- and L-moves. The point estimate in this column therefore indicates that the frequency of mistakes decreases, on average, by .279 p.p. for every W-move that is added to the choice set. This corresponds to almost 5% of the mean frequency of mistakes. In analogous fashion, the coefficients in cols. (2) and (3) imply that mistake frequencies respectively increase by .810 and .247 p.p. as D- and L-moves are added to the choice set. The evidence in the first three columns of Table 5 thus shows that whether mistakes become more or less frequent as the choice set expands depends on the type of move that is being added. This pattern helps to explain why the raw data yield little evidence of choice overload, i.e., why mistake frequencies are not uniformly higher for larger choice sets. Perhaps more importantly, the increase in the frequency of mistakes is significantly greater when a D- rather than an L-move is added to the choice set (p < .001). This finding is consistent with Prediction 4.

 $^{^{24}}$ In Appendix Table AT.2, we show that controlling for object complexity is not necessary to obtain results that are substantively equivalent to those in Table 5.

	TAULT	the second secon		rendrinn apripar			
	(1)	(2)	(3)	(4)	(5)	(9)	(2)
			ц	robability of Mista	ke		
Number of W-Moves $(\div 100)$	-0.279 (0.001)			$^{-1.242}(0.002)$	-0.988 (0.003)		
Number of <i>D</i> -Moves $(\div 100)$ Number of <i>L</i> -Moves $(\div 100)$		0.810 (0.002)	0.247			0.841 (0.003)	
Total Number of Moves $(\div 100)$			(0.003)				-0.017 (0.001)
Fixed Effects:							
Number of D - × L -Moves	Yes	No	No	No	No	No	No
Number of $W - \times L$ -Moves	No	Yes	No	No	No	No	No
Number of $W - \times D$ -Moves	No	No	Yes	No	No	No	No
Total Moves \times Number of L-Moves	No	No	No	Yes	No	No	No
Total Moves \times Number of <i>D</i> -Moves	No	No	No	No	Yes	No	No
Total Moves \times Number of W-Moves	No	No	No	No	No	Yes	No
Fraction of W - $\times D$ - $\times L$ -Moves	No	No	No	No	No	No	Yes
Player	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Minimal DTM among W -Moves	Yes	\mathbf{Yes}	Yes	\mathbf{Yes}	Yes	Yes	Yes
R^2	0.119	0.144	0.153	0.120	0.143	0.151	0.155
N	226,955,095	226,955,095	226,955,095	226,955,095	226,955,095	226,955,095	226,955,095
Notes: Entries are coefficients and standar the minimal DTM among W -moves in the a decision, and individual observations are All estimates are scaled to correspond to t errors are two-way clustered by player and	d errors from estin same choice set. C : reweighted so tha the percentage poin l game, and are sho	anating η in variant: ontrols for the num t all players receive t change in the pro- own in parentheses	s of eq. (6) by ordin aber and types of o e equal total weight obability of an erro	ary least squares. ther moves vary acr The sample in ev r associated with a	All regressions contracts contracts columns. The u ery column includes one-unit increase ir	rol for player fixed nit of observation i s all board configur the respective reg	effects as well as n each regression is ations in our data. ressor. Standard

Table 5: Mistakes as a Function of Choice-Set Composition

Changes to the choice set that involve replacing one move with another one of a different type are considered in cols. (4)–(6). The coefficient in col. (4) refers to a W-move being replaced by a D-alternative, whereas that in col. (5) refers to a switch between W- and L-moves. The point estimate in col. (6) tells us that, on average, mistake frequencies increase by .841 p.p.—or about 15% of the mean mistake frequency—when an L-move is replaced by a D-alternative.

Col. (7) considers homothetic enlargements of the choice set, i.e., expansions that hold the *share* of W-, D-, and L-moves fixed. After carefully controlling for the composition of available alternatives, we find no evidence to suggest that larger choice sets are associated with more mistakes. Not only is the coefficient in col. (7) more than an order of magnitude smaller than the other estimates in Table 5, it is, if anything, negative and very precisely estimated.

In sum, the results in Table 5 are consistent with the theory's predictions. Moreover, the point estimates in cols. (1)–(6) are nontrivial in size—ranging from 4% to 21% of the mean mistake rate in the data. Our findings, therefore, imply that complexity at the set level is empirically relevant and closely related to the composition rather than the size of the choice set.

7.3. Set-Level Complexity Beyond Chess

Motivated by these findings, we propose a concise, portable definition of what we call composition complexity. Let A be a choice set, with $V(A) = \{v_x \mid x \in A\}$ denoting the set of values in A. Let $\bar{v}(A)$ be the largest value in V(A), and let N(v, A) be the number of objects in A that have value v. With this notation in hand and assuming that |V(A)| > 1, the composition complexity of A is:

(7)
$$\Gamma(A) \equiv -N(\bar{v}(A), A) + \sum_{v \in V(A) \setminus \bar{v}(A)} \alpha(\bar{v}(A) - v) N(v, A),$$

where $\alpha(\cdot)$ is a positive and monotonically decreasing function.

The first term in this expression implies that as the number of maximal objects in the choice set increases, the choice task becomes easier. The function $\alpha(\cdot)$ in the second term captures the key features of the analysis above. Positivity of $\alpha(\cdot)$ means that adding non-maximal objects increases complexity, and monotonicity implies that complexity increases as suboptimal objects become more similar to optimal ones.²⁵

²⁵The exact cardinal specification of $\alpha(\cdot)$ determines, together with the composition of the choice set, how composition complexity changes when an object with value above $\bar{v}(A)$ is added to A.

8. Concluding Remarks

This paper studies how complexity at the level of individual objects affects decision making. We establish that object complexity affects choice in systematic ways. Choice frequencies of optimal alternatives decrease in the respective object's complexity, whereas choice frequencies of suboptimal alternatives follow the opposite pattern. As a consequence, the frequency of mistakes increases (decreases) as any of the available (sub)optimal alternatives becomes more complex. We also develop an empirical test that leverages complexity to distinguish between satisficing and a large class of maximization-based models. Based on this test, we reject maximization as the primary choice procedure in our data.

Taking a step back, our analysis hints at the potential of developing a ranking of choice sets according to their complexity. Take two sets A and B with the same support of values. Based on our results, set B is more complex than set A if it is possible to transition from Ato B by iteratively applying any of the following operations:

- increasing the object complexity of some alternative,
- adding a suboptimal alternative, or
- removing an optimal one.

The resulting binary relation is transitive but incomplete, suggesting an avenue for future research.

More broadly, our results on complexity in decision-making offer insights that may be important in other real-world settings. For example, if consumers' choice sets include products by different firms, then firms with superior offerings may wish to simplify their product descriptions. In contrast, firms with inferior offerings have an incentive to make their products more complex. Firms also have an incentive to affect the order in which satisficing consumers consider offerings even if doing so is costly. Similar incentives may be present in the design and presentation of policy proposals, in strategic games, and in various other settings.

References

- Abeler, J. and S. Jäger (2015). Complex Tax Incentives. American Economic Journal: Economic Policy 7(3), 1–28.
- Abreu, D. and A. Rubinstein (1988). The Structure of Nash Equilibrium in Repeated Games with Finite Automata. *Econometrica* 56(6), 1259–1281.
- Anderson, A., J. Kleinberg, and S. Mullainathan (2017). Assessing Human Error Against a Benchmark of Perfection. ACM Transactions on Knowledge Discovery from Data (TKDD) 11(4), 45.
- Bossaerts, P. and C. Murawski (2017). Computational Complexity and Human Decision-Making. Trends in Cognitive Sciences 21(12), 917–929.
- Camerer, C. (2003). Behavioral Game Theory. Princeton, NJ: Princeton University Press.

- Caplin, A., M. Dean, and D. Martin (2011). Search and Satisficing. American Economic Review 101(7), 2899–2922.
- Casella, G. and R. Berger (2001). Statistical Inference (2nd ed.). Boston: Cenage.
- Chernev, A., U. Böckenholt, and J. Goodman (2015). Choice Overload: A Conceptual Review and Meta-Analysis. *Journal of Consumer Psychology* 25(2), 333–358.
- Chiappori, P.-A., S. Levitt, and T. Groseclose (2002). Testing Mixed-Strategy Equilibria When Players Are Heterogeneous: The Case of Penalty Kicks in Soccer. American Economic Review 92(4), 1138–1151.
- Enke, B. and T. Graeber (2021). Cognitive Uncertainty. NBER Working Paper No. 26518.
- Gabaix, X., D. Laibson, G. Moloche, and S. Weinberg (2006). Costly Information Acquisition: Experimental Analysis of a Boundedly Rational Model. *American Economic Review 96*(4), 1043–1068.
- Glickman, M. E. (1999). Parameter Estimation in Large Dynamic Paired Comparison Experiments. Journal of the Royal Statistical Society: Series C 48(3), 377–394.
- Glickman, M. E. (2000). The Glicko-2 System for Rating Players in Head-to-Head Competition. mimeographed, Boston University.
- Hsu, S.-H., C.-Y. Huang, and C.-T. Tang (2007). Minimax Play at Wimbledon: Comment. American Economic Review 97(1), 517–523.
- Huck, S. and G. Weizsäcker (1999). Risk, Complexity, and Deviations from Expected-Value Maximization: Results of a Lottery Choice Experiment. Journal of Economic Psychology 20(6), 699–715.
- Iyengar, S. and M. Lepper (2000). When Choice is Demotivating: Can One Desire Too Much of a Good Thing? Journal of Personality and Social Psychology 79(6), 995–1006.
- Jakobsen, A. M. (2020). A Model of Complex Contracts. American Economic Review 110(5), 1243–1273.
- Kalai, E. and W. Stanford (1988). Finite Rationality and Interpersonal Complexity in Repeated Games. *Econometrica* 56(2), 397–410.
- Levitt, S. and J. List (2007). What Do Laboratory Experiments Measuring Social Preferences Reveal About the Real World? *Journal of Economic Perspectives* 21(2), 153–174.
- Levitt, S., J. List, and D. Reiley (2010). What Happens in the Field Stays in the Field: Exploring Whether Professionals Play Minimax in Laboratory Experiments. *Econometrica* 78(4), 1413–1434.
- Levitt, S., J. List, and S. Sadoff (2011). Checkmate: Exploring Backward Induction among Chess Players. *American Economic Review* 101(2), 975–990.
- Luce, D. (1959). Individual Choice Behavior. New York: Wiley.
- Man, R. d. (2013). Syzygy tablebases. available at https://github.com/syzygy1/tb.
- Marschak, J. (1960). Binary-Choice Constraints and Random Utility Indicators. In K. Arrow,S. Karlin, and P. Suppes (Eds.), Stanford Symposium on Mathematical Methods in the Social

Sciences, pp. 312–329. Stanford, CA: Stanford University Press.

- McFadden, D. (1974). Conditional Logit Analysis of Qualitative Choice Behavior. In P. Zambreka (Ed.), *Frontiers in Econometrics*, pp. 105–142. New York: Academic Press.
- Nalimov, E., G. Haworth, and E. Heinz (2000). Space-Efficient Indexing of Chess Endgame Tables. *ICGA Journal* 23(3), 148–162.
- Neumann, J. v. (1928). Zur Theorie der Gesellschaftsspiele. Mathematische Annalen 100(1), 295–320.
- Neumann, J. v. and O. Morgenstern (1944). Theory of Games and Economic Behavior. Princeton, NJ: Princeton University Press.
- Neyman, A. (1985). Bounded Complexity Justifies Cooperation in the Finitely Repeated Prisoners' Dilemma. *Economics Letters* 19(3), 227–229.
- Oprea, R. (2020). What Makes a Rule Complex? American Economic Review 110(12), 3913–3951.
- Osborne, M. and A. Rubinstein (1994). A Course in Game Theory. Cambridge, MA: MIT Press.

Palacios-Huerta, I. (2003). Professionals Play Minimax. Review of Economic Studies 70(2), 395–415.

- Palacios-Huerta, I. and O. Volij (2008). Experientia Docet: Professionals Play Minimax in Laboratory Experiments. *Econometrica* 76(1), 71–115.
- Palacios-Huerta, I. and O. Volij (2009). Field Centipedes. American Economic Review 99(4), 1619–1635.
- Rubinstein, A. (1986). Finite Automata Play the Repeated Prisoner's Dilemma. Journal of Economic Theory 39(1), 83–96.
- Rubinstein, A. (2007). Instinctive and Cognitive Reasoning: A Study of Response Times. *Economic Journal* 117(523), 1243–1259.
- Rubinstein, A. (2016). A Typology of Players: Between Instinctive and Contemplative. Quarterly Journal of Economics 131(2), 859–890.
- Salant, Y. (2011). Procedural Analysis of Choice Rules with Applications to Bounded Rationality. American Economic Review 101(2), 724–748.
- Schwartz, B. (2004). The Paradox of Choice: Why More Is Less. New York: HarperCollins.
- Simon, H. (1955). A Behavioral Model of Rational Choice. Quarterly Journal of Economics 69(1), 99–118.
- Simon, H. (1972). Theories of Bounded Rationality. In C. McGuire and R. Radner (Eds.), Decision and Organization, pp. 161–176. Amsterdam: North-Holland.
- Spenkuch, J., P. Montagnes, and D. Magleby (2018). Backward Induction in the Wild? Evidence from Sequential Voting in the US Senate. American Economic Review 108(7), 1971–2013.
- Thompson, K. (1986). Retrograde Analysis of Certain Endgames. ICGA Journal 9(3), 131–139.
- Walker, M. and J. Wooders (2001). Minimax Play at Wimbledon. American Economic Review 91(5), 1521–1538.
- Wilson, A. (2014). Bounded Memory and Biases in Information Processing. *Econometrica* 82(6), 2257–2294.
- Wooders, J. (2010). Does Experience Teach? Professionals and Minimax Play in the Lab. Economet-

rica 78(3), 1143-1154.

- Woodford, M. (2020). Modeling Imprecision in Perception, Valuation, and Choice. Annual Review of Economics 12, 579–601.
- Zermelo, E. (1913). Über eine Anwendung der Mengenlehre auf die Theorie des Schachspiels. In E. Hobson and A. Love (Eds.), Proceedings of the Fifth International Congress of Mathematicians (Cambridge 1912), pp. 501–504. Cambridge, UK: Cambridge University Press.