

RESEARCH ARTICLE

SMS | Strategic Management Journal

WILEY

Training with AI: Evidence from chess computers

Fabian Gaessler^{1,2}  | Henning Piezunka³ 

¹Pompeu Fabra University, Barcelona, Spain

²Max Planck Institute for Innovation and Competition, Munich, Germany

³INSEAD, Fontainebleau, France

Correspondence

Henning Piezunka, INSEAD, Fontainebleau, France.

Email: henning.piezunka@insead.edu

Abstract

We suggest that AI can help decision-makers learn; specifically, that it can help them learn strategic interactions by serving as *artificial* training partners and thus help them to overcome a bottleneck of scarce human training partners. We present evidence from chess computers, the first widespread incarnation of AI. Leveraging the staggered diffusion of chess computers, we find that they did indeed help chess players improve by serving as a substitute for scarce human training partners. We also illustrate that chess computers were *not* a perfect substitute, as players training with them were not exposed to and thus did not learn to exploit idiosyncratic (“human”) mistakes. We discuss implications for research on learning, on AI in management and strategy, and on competitive advantage.

KEYWORDS

artificial intelligence, chess, difference-in-differences, learning, strategic interaction

1 | INTRODUCTION

We suggest that AI can help decision-makers learn. Specifically, we propose that AI-backed computer simulations can train decision-makers in strategic interactions. While strategic

Fabian Gaessler and Henning Piezunka have contributed equally to this study.

This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial-NoDerivs](https://creativecommons.org/licenses/by-nc-nd/4.0/) License, which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

© 2023 The Authors. *Strategic Management Journal* published by John Wiley & Sons Ltd.



interactions are ubiquitous (e.g., when competing or negotiating), learning them is challenging because gathering experience is expensive. Decision-makers therefore often choose to train strategic interaction with training partners who emulate actual opponents. But qualified training partners are often scarce—a bottleneck for decision-makers learning strategic interactions.

We propose that AI can help overcome that bottleneck. Specifically, we suggest that AI-backed computer simulations provide decision-makers with *artificial training partners* that can realistically emulate the behavior of actual opponents. Thanks to continuous improvements in algorithms and computational power, AI-backed computer simulations are a scalable substitute for scarce human training partners and thus help decision-makers learn strategic interactions. However, they are not necessarily a perfect substitute. When training with boundedly rational human training partners, decision-makers are exposed to “human” mistakes which consistently performing AI does not make. Decision-makers who train with AI are thus not exposed to (human) mistakes and may—as a consequence—*not* learn to recognize and exploit them.

We examine the potential of AI to help decision-makers learn strategic interaction by analyzing the effect of chess computers on the performance of chess players. This empirical setting allows us to link the first widespread application of AI (chess computers) with a prime example of complex strategic interaction (chess). We use data on more than 20,000 chess players in Western Europe and the former Soviet Union (hereafter, just “Soviet Union”) and analyze their performance across half a million tournament games. Our window of observation is 1970 to 2000, a period in which chess computers diffused widely and evolved tremendously. We link this data with detailed historical information on commercial chess computers. While chess, being a *well-defined problem* in a closed setting, is unrepresentative of many real-life strategic challenges, it does constitute a *complex strategic interaction* and thus serves as an early testbed to examine whether decision-makers can train such interactions with the help of AI.

We start by investigating the *main effect*; that is, whether chess computers have a positive effect on players’ performance. In a difference-in-differences (DiD) framework, we leverage two natural experiments that offer exogenous albeit high-level variation in *access to chess computers*: first, the introduction of commercial chess computers in Western Europe, beginning in 1977, and, second, their belated diffusion in the Soviet Union after the fall of the Iron Curtain in 1989. We find an increase in chess performance among Western European players from 1977 and a relative catch-up of Soviet players from 1989. We provide additional evidence that the performance increase is indeed due to chess computers. Most notably, we exploit a particular limitation of early chess computers that made them more effective opponents when they played black. Our finding that players with access to chess computers improved more in playing white strongly suggests that it was chess computers that helped them improve.

We also illustrate an important boundary condition: the effect of chess computers is confined to players whose skills are inferior to the chess computers available at a given time. Given the tremendous evolution of strength in chess computers, this boundary condition implies that weak players tended to “catch up,” resulting in a convergence of the skill distribution.

We test the suggested *mechanism* of how AI helps in learning strategic interaction by being a substitute for scarce human training partners. We distinguish players by how many opportunities they have to train with human training partners. We find that players with *fewer* such opportunities benefit *more* from chess computers. By extension, our argument suggests that players with access to chess computers may forgo participation in tournaments that serve primarily as an opportunity to train with other humans. Our data provide consistent evidence: players with access to chess computers play fewer tournaments.

The finding that artificial training partners are a substitute for human training partners in learning strategic interaction raises the question of whether they are a *perfect substitute*. We therefore examine *whether the learning outcome differs between training with AI and training with a human*. We focus on a distinct feature that separates AI from humans: the consistency of its performance. Boundedly rational human training partners make idiosyncratic mistakes—so-called blunders—due to lack of concentration, but AI does not; “machines don’t get tired.” Decision-makers who train with AI thus have less exposure to blunders, but part of succeeding in competitive strategic interaction is to recognize and exploit an opponent’s blunder. We indeed find players with access to chess computers less able to exploit a human opponent’s blunder.

Our findings allow us to contribute to three research domains. By examining how decision-makers learn strategic interactions, we contribute to research on *learning* (Anand et al., 2016; KC et al., 2013; Maula et al., 2022). We illustrate that scarce human training partners can be a bottleneck in learning strategic interactions and that AI can help overcome this bottleneck, as it is a scalable substitute. That said, AI is not necessarily a perfect substitute, as the outcomes from training with human and artificial training partners differ.

We also contribute to research on the role of AI in management and strategy. While research has pointed out that AI can either substitute for or complement human decision-makers (Agrawal et al., 2019; Brynjolfsson & Mitchell, 2017; Felten et al., 2021), we show that it can also help them *learn*.

Our findings inform research on the dynamics of competitive advantage. We complement research on population-level performance distribution (Kapoor & Agarwal, 2017; Lenox et al., 2006) by illustrating how new training technologies can change it. We contribute to research on strategic interaction by illustrating that the ability to exploit the blunders of boundedly rational opponents is an important skill for a decision-maker and is better acquired through training with boundedly rational humans than with consistently high-performing machines.

2 | THEORETICAL BACKGROUND

2.1 | The difficulty of learning strategic interactions

It is crucial for decision-makers to learn strategic interactions. Making mistakes in strategic interactions can be very costly as they are often irreversible and highly consequential (e.g., a company may slide in the competitive rank order). It is thus crucial that organizations and their decision-makers find a way to learn strategic interactions.

It is, however, challenging to learn strategic interactions as its process differs from that of learning other kinds of activities. The classical learning process is composed of the actor taking some action and receiving performance feedback (Greve, 2003; Joseph et al., 2016; Keil et al., 2022), allowing the trainee to learn about the effectiveness of an action. In strategic interactions, a sequence of actions and *responses* takes place—and only at the end is the actual outcome revealed.¹ Learning strategic interactions thus requires learning a particular kind of skill:

¹For example, prior literature has considered location choices and alliance building as strategic interactions at the firm level (Alcácer et al., 2015; Panico, 2017). More generally, any bargaining between firms (e.g., auctions, litigation) or within firms (e.g., salary negotiations) can be understood as a strategic interaction.



anticipating the opponent's response (Levine et al., 2017; Menon & Yao, 2017). Beyond requiring an additional kind of skill, the outlined sequence of action and responses also accentuates many learning challenges. For example, the sequence of actions and responses results in numerous potential scenarios, making it hard to gain an overview of the performance landscape (Adner et al., 2014; Siggelkow, 2002). Also, the effect of any action is contingent on the actions taken by other actors—making learning harder (Piezunka et al., 2022; Rivkin & Siggelkow, 2003). Moreover, while credit assignment is generally difficult (Clough & Piezunka, 2020; Rahmandad et al., 2009), this is particularly true when performance is revealed only at the end (Christensen & Knudsen, 2010; Fang & Levinthal, 2009)—like it is often the case in strategic interactions. In brief, the interplay of actions and responses renders strategic interactions difficult to learn.

2.2 | The need for and benefits of experiential learning

The features that render strategic interactions difficult to learn also limit the progress that can be made in a *non-experiential manner*. In non-experiential learning, decision-makers rely on codified knowledge (e.g., books or lectures). They thus learn about potential actions and responses but do not have the actual experience of taking such actions and receiving such responses (Martin et al., 2014). Learning from codified knowledge can provide a general understanding and a familiarity with archetypical strategies, but is unlikely to be sufficient to master strategic interactions. Even if the codification is extensive, the complexity of strategic interactions will unavoidably result in non-codified scenarios, where the actions that the trainee learned to be effective turn out to be ineffective. Experiential learning is thus crucial as it can accustom decision-makers to navigating unknown scenarios. Learning from codified information is also insufficient insofar as strategic interaction requires not only understanding but also implementation. All told, the challenges in learning strategic interactions and the insufficiencies of learning them non-experientially suggest that experiential learning is essential.

Learning strategic interaction experientially implies taking an action *and* receiving a response. This underscores a crucial difference between learning strategies interactions and many other activities—the need for an interaction partner who responds. Receiving a response helps decision-makers learn about the effectiveness of their actions; it allows them to become better in the crucial task of anticipating an opponent's response (Levine et al., 2017). Decision-makers can also explore alternative actions, gaining an understanding of the actions' relative effectiveness and extending their own repertoire of actions (Smith et al., 2005). Experiential learning also allows them to learn about non-codified (and typically less common) scenarios (Martignoni et al., 2020). Taken together, experiential learning, with decision-makers receiving responses to their actions, allows them to improve in strategic interactions.

2.3 | Traditional ways of learning strategic interactions experientially

One way to learn experientially is *to engage in actual strategic interactions*; that is, to learn “online” (Gavetti & Levinthal, 2000). However, learning this way is challenging. First, opportunities to engage in strategic actions are often limited. Second, the stakes in strategic actions are often high (relative to other organizational actions), rendering suboptimal outcomes expensive. Third, the actions taken are often irreversible and may cast a long shadow by affecting

subsequent actions and responses as well. Finally, there may be undesired externalities, such as that potential strategies to competitors are revealed, which reduces the chance to surprise them with a well-proven action. In brief, due to their sequential and repeated nature, learning strategic interactions through actual engagement is very costly.

Given the benefits of experiential learning and the costs of engaging in actual strategic interactions, decision-makers often learn strategic interactions by simulating them with human training partners. Executives may hire a consulting company to engage in war games, with the consultants emulating potential strategic responses to strategic actions the firm is considering (Courtney et al., 2020). But while human training partners have the benefit of being “responsive,” they are often scarce. This is because responsive training is usually bilateral (Bloom, 1984); that is, training with a human partner is hard to scale. Human training partners may also be hard to find and/or costly. This is particularly true for human training partners sufficiently skilled to emulate the strategic responses of actual opponents. All told, training with human training partners is effective, but they are also scarce and thus constitute a potential bottleneck in learning strategic interaction.

2.4 | Learning strategic interactions experientially with AI

We suggest that AI can serve as a *substitute* for human training partners. We conceptualize the AI to be part of a computer-based simulation, with the latter providing the environment for the strategic interaction. The role of the AI in such simulations is to emulate competitors and their responses to the decision-maker's actions—and thus serve as *artificial training partners*.

We are *not* the first to point out that people can use technology for training (Bailey et al., 2019; Beane, 2019). People already train various activities with the help of *conventional simulations* (Salas et al., 2012). For example, tennis-ball machines help tennis players train by serving balls, and flight simulators help pilots familiarize themselves with air navigation. While conventional simulations sometimes involve interaction partners, they follow simple routines and “if-then” expressions, which do not reflect an actual human decision-making process. The responses of conventional simulations thus often lack the necessary realism to make them meaningful training partners, especially in complex strategic interactions. In short, conventional simulations are helpful for training certain activities, but their unrealistic responses make them fall short in teaching complex strategic interactions.

AI-backed simulations, however, allow for more realistic emulation of competitors than previously possible due to advances in predictive computer algorithms—such as neural networks—combined with ever-increasing computational power and data availability (Agrawal et al., 2018; Wu & Brynjolfsson, 2015). We therefore suggest that AI-backed simulations can be artificial training partners that emulate competitors in their responses.

An example provided to us by a military training expert illustrates the difference that AI can make. Consider a simulation in which soldiers have to conquer a territory. In a conventional simulation without AI, the emulated enemy soldiers each react to the trainees' actions according to relatively simple and predefined—thus, not very realistic—routines. With AI, the emulated enemies coordinate among one another and respond to the soldiers' actions as real enemy soldiers would. This difference can be crucial for learning as it separates a toy-like simulation from a sufficiently realistic training simulation. That said, the ability of AI to emulate human behavior can vary substantially depending on the underlying algorithm and the specific task (Athey et al., 2020; Dell'Acqua, 2022).




FIGURE 1 Schematic overview of methods to train strategic interactions.

The particular appeal of AI-backed simulation as a mode to learn strategic interactions experientially is thus twofold: (a) AI-backed simulations are scalable and thus highly available (as are books, lectures, and conventional computer-based simulations), and (b) AI-backed simulations provide realistic responses (as human training partners do). When one maps out the training modes discussed so far—as in Figure 1—the special position and particular appeal of AI-backed simulation become evident: *it resolves the tradeoff between availability and responsiveness* that has so far defined the landscape of training modes.

2.5 | Research questions

The arguments above raise several research questions, which we examine empirically.

First: *Do decision-makers benefit from training with AI-backed computer simulations?* That is, does their performance in actual strategic interactions improve? We seek to explore this question “in the field,” where other modes of training strategic interactions already exist. In other words, we do *not* ask whether training with AI-backed simulations is better than no training but whether it is beneficial in the presence of other training modes.

Second: Is the proposed effect of AI-backed simulations subject to a *boundary condition*? We suggest that artificial training partners provide the kinds of responses that decision-makers would receive from *actual* competitors. It therefore seems plausible that AI-backed simulations only benefit decision-makers if the AI can emulate the responses of their actual competitors.

Third: Is *substitution* for human training partners the core *mechanism*? In other words, do AI-backed simulations help decision-makers learn strategic interaction because artificial training partners solve the bottleneck of scarce human training partners by being a scalable substitute?

Fourth: Are AI-backed simulations a *perfect* substitute for human training partners? That is, do the learning outcomes with AI-backed simulations match those with human training partners?

3 | SETTING AND DATA: CHESS COMPUTERS

To examine whether AI can help people improve at strategic interactions, we study the effect of chess computers on the performance of (human) chess players.

3.1 | Chess as Strategic Interaction, Chess Computers as AI-backed Simulations

We use chess to examine decision-makers' ability to learn strategic interactions. In chess, two individuals compete in a closed setting, each taking an action followed by the opponent's response. Given the interactive nature of the game, the large number of scenarios, and the difficulty of credit assignment, chess is widely considered a prototypical—and is an often-cited—example of strategic interaction (Krakowski et al., 2022).²

To examine whether and how AI helps decision-makers learn strategic interactions, we study the effect of *chess computers*. The link between AI and chess computers goes back to AI's inception. Alan Turing, the father of modern computer science, and Claude Shannon, the father of modern information theory, both studied chess. Since then, chess has been the primary use case for AI innovation, and chess computers have been at the frontier of AI applications (Kasparov, 2018). Herbert Simon, a pioneer of AI and organization theory, described chess as the fruit fly of AI—that is, a simple system that allows the initial exploration of complex phenomena (Chase & Simon, 1973), as the fruit fly often does in the study of human genetics. As “machine[s] [able] to imitate intelligent human behavior” (Merriam Webster, 2022), chess computers became, in the 1970s, the first commercial application of an AI-backed simulation.

Although there has not been a lasting definition of AI,³ chess computers exemplify two of its core aspects. First, AI has generally been associated with *complex* tasks, such as translation, hiring, and medical diagnosis (Agrawal et al., 2018; Brynjolfsson et al., 2019; Felten et al., 2021). Chess is notoriously complex; there are 69,352,859,712,417 possible scenarios after only the first 10 moves. The selective evaluation of these scenarios challenges the cognitive abilities of even the most intelligent humans. Second, AI has been associated with *emulating human thought* (Simon, 1996). A human chess player narrows the search scope based on some (implicit) rules and then evaluates only certain moves in depth. The AI in chess computers is programmed to apply such selective evaluation processes by following a predetermined set of heuristics (Simon & Schaeffer, 1992). In contrast, modern AI relies primarily on machine learning techniques in which the AI follows rules derived from existing data. Despite this difference in the AI's backend, its defining purpose—emulating human decision-making—has remained the same.

It is important to note that the typical chess computer in our sample only responded to the focal player's moves with a countermove; it did not provide additional feedback (e.g., an explanation of why the player's move was not optimal). In this regard, chess computers are probably best compared with training partners that take on the role of real opponents rather than to teachers or coaches; that is, they respond to the player's move but provide no further feedback.

²More broadly, a rich emerging body of research has illustrated the potential of leveraging sports data for research in management and competition in particular (Marino et al., 2015; Sharapov & Ross, 2023)

³Definitions of AI vary mostly in their breadth: artificial intelligence is usually seen as an umbrella term for all machines that imitate human behavior/intelligence. *Machine learning* methods are a subset of AI methods and *deep learning* methods are a subset of machine learning. The debate on what defines AI is exemplified by the so-called *AI effect* (McCorduck, 2004)—the repeated observation that as soon as AI solves a problem, that problem is no longer seen as requiring actual intelligence.



3.2 | Data

Our research design combines data on chess players and chess computers. We create a three-decade-long unbalanced panel of about 20,000 *players* with detailed information on more than 500,000 tournament games. We extract player- and game-level data from ChessBase, the leading chess database providing comprehensive coverage of tournament games. ChessBase includes meta-information on more than seven million games—such as date, location, tournament, and outcome—and on the players—such as name, country, and birth year. We transform the dataset from game to player-year level. We focus on 1970 to 2000—the period in which chess computers became available and diffused widely. Our player sample is defined by tournament activity and country of residence. All players in our sample have played tournament games. We restrict our sample to players residing in either Western Europe or the Soviet Union. ChessBase includes the moves played in each game; we process this granular information on more than 21 million moves to capture the players' performance in strategic interactions.⁴

We also create a dataset on chess computers. We first assemble an exhaustive list of over 500 commercial models with information on the release date, strength, price, and manufacturer.⁵ We then determine the upper-bound strength of all available chess computers for each year from their introduction in 1977 to 2000.

4 | ANALYSIS

4.1 | Analysis—Part 1: Main effect

Our analysis starts out by examining *whether* chess computers have a positive effect on players' performance. We leverage two natural experiments in the form of the staggered diffusion of chess computers in Western Europe and in the Soviet Union.

4.1.1 | Research design

We use the staggered diffusion of chess computers as an exogenous variation in *access to chess computers*. Chess computers became available in Western Europe in 1977,⁶ but not in the Soviet

⁴The overwhelming majority of players in our data reside in Western Europe (see Supporting information Figure A1). Very few are active for the full period; most either drop out before the panel ends (retirees) or join after it begins (newcomers). In general, due to the increasing popularity and coverage of chess, the data become denser (more players, more games per player) towards the end of the panel. We address the panel's unbalancedness in robustness checks.

⁵Our primary source for this is <http://www.schach-computer.info/wiki>, a community-based wiki. We cross-check and supplement this source with information from various online and offline sources.

⁶In 1977, the first commercial chess computer, *Chess Challenger*, went on sale in the West and quickly diffused. Several manufacturers introduced stronger and cheaper models over the following years (see Supporting information Figure A2). See Figure A3 for the typical design of a dedicated chess computer. Although systematic sales data are not available for this period, we know that (a) worldwide sales were about 300,000 in 1979 and 1.4 million in 1981; (b) in 1987, more than 1 million West Germans had access to a chess computer and 1.8 million intended to buy one; and (c) in 1989, the global chess computer market was valued at about 45 million USD annually. These figures are in line with anecdotal evidence of the widespread adoption of chess computers in Western Europe. For instance, in 1982, the US Chess Federation promoted a branded chess computer exclusively for members (see Figure A4). In the same year, computer chess companies first topped 100 million USD in sales.

Union until 1989.⁷ These two discrete events allow us to run a DiD analysis with the following pattern: From 1970 to 1976, no player had access to chess computers. From 1977 to 1989, only players in Western Europe had access. From 1989 onward, players in Western Europe and the Soviet Union had access. This pattern should imply that any estimated differential effect of chess computers on player performance (that is, in Western Europe vs. the Soviet Union) should vanish once they became available in both regions.

4.1.2 | Variables and econometric model

Dependent variables: To examine whether access to chess computers helps players learn strategic interactions, we measure players' progress in performance. We construct three distinct measures of player performance. Our primary measure—*Elo rating*—is used to quantify the relative skill of players in many zero-sum games, such as chess. Elo rating increases (decreases) if the player wins (loses), with the increase (decrease) dependent on the difference in skill level between the two players. Elo rating is well established in chess and is directly provided by our secondary data. It may, however, be subject to biases.⁸

We therefore use additional performance measures. One, *average centipawn loss*, is based on centipawn loss, a common measure of the quality of a single chess move.⁹ It is determined by comparing a player's advantage from playing a chosen move with the advantage from playing the best move available. To use this measure, we assessed all moves in our sample (21 million) and identified the optimal move in each case. To this end, we use Stockfish, a high-end open-source chess engine.¹⁰ We then determine the player's advantage before and after the chosen move and calculate the loss or gain in advantage relative to that of the best available move (see Supporting information Appendix C for details). The *average centipawn loss* is the player-year-specific average difference in

⁷Chess computers were de facto unavailable in the Soviet Union. In fact, the Soviet Computer Chess Federation noted that “[w]ith the intense interest in chess of the Soviet people and the early success of computer chess, it is perhaps surprising that from the mid-70s to 1988 there was little computer-chess activity in the Soviet Union” (Donskoy & Schaeffer, 1988). No competitive product was developed or produced in the Soviet Union or its satellite states and the 1979 Western trade embargo on electronics (the COCOM embargo) made it almost impossible to import microchip-based devices (Gustafson, 1981). We cannot fully exclude the availability of chess computer software, as there was a lot of software plagiarism (circumventing the trade embargo). However, computers remained a luxury item for most of the Soviet population (Judy & Clough, 1990). Only with Glasnost (1987) and then the fall of the Iron Curtain did the import of Western electronics start to increase. The first Soviet private chess computer clubs were founded in 1989 and the leading Soviet chess magazine established a regular chess computer column in 1990.

⁸The Elo rating is comparative and therefore depends on the pool of players. In our empirical analysis, we compare the performance of Western and Soviet players with each other. Differences in pool composition (e.g., a larger share of bad players in tournaments in one region) could render their Elo ratings less comparable. We compare the distribution of “home” and “away” games of Western and Soviet players before 1989 (see Supporting information Figure A6) and find that despite the Iron Curtain, Soviet and Western players frequently encountered each other in international tournaments. We therefore gain confidence that we can take their Elo ratings at face value.

⁹The centipawn is a chess-internal unit of measure to quantify advantage. A centipawn equals 100th of a pawn (the least-valued piece in chess). Centipawns play no formal role in the game but prove helpful in evaluating positions and comparing possible moves.

¹⁰Stockfish has been consistently ranked first or second in chess-engine rating lists and is widely seen as one of the strongest conventional chess engines in the world. For more information, see <https://stockfishchess.org/>. Stockfish searched 20 moves deep and evaluated up to 70 million possible moves per second. Given the large number of moves to be evaluated, we faced a parallel runtime of about 6 weeks with 96 virtual central processing units on a cloud computing platform.



advantage from all played moves. As we illustrate in Supporting information Figure A5, Elo rating and average centipawn loss are highly correlated: a player with a high (low) Elo rating has a small (large) average centipawn loss. Thus, if access to chess computers has a positive effect on chess performance, it should have a *negative* effect on the average centipawn loss.

For our third measure of performance, *Game won*, we observe the outcome of a given game from the player's perspective, focusing exclusively on direct encounters between Western and Soviet players, who frequently face off at tournaments despite political boundaries.

Independent variables: Our key independent variable, *Chess computer access*, captures the possibility of accessing chess computers. It is time-variant and indicates whether chess computers were available in a player's region. Western players have had access to chess computers since 1977, but Soviet players only since 1989. *Chess computer access* is therefore a binary variable that equals 0 for all players starting in 1970, equals 1 for Western players from 1977 on, and equals 1 for Soviet players from 1989 on. Our treatment variable captures whether the focal player resided in a region with access to chess computers, not whether the player actually used one. This limitation most likely results in an underestimation of the actual effect.

Econometric model: For the main part of our empirical analysis, we rely on DiD models with player fixed effects, exploiting the treatment condition (access to chess computers) over time for the effect of chess computers on player performance. We create a panel at the player-year level and compare the performance of player i over time t to capture learning, distinguishing years in which i is either treated or not treated. The model can be written:

$$y_{it} = \alpha_i + \beta_1 \text{Chess computer access}_{it} + \text{age}_{it} + \delta_t, \quad (1)$$

where y is the dependent variable, capturing a player's performance. Since performance may correlate with general advances in chess theory and didactic, we control for time trends using calendar-year fixed effects. Because treatment might be correlated with the player's seniority, we include age dummies to capture nonlinear life-cycle effects. Identification relies on the comparability of players from Western Europe and the Soviet Union. We discuss the latter's virtue as a control group below and solidify the common trend assumption through an event study. In all models, we cluster standard errors at the player level.

4.1.3 | Results

We find a *positive effect* of access to chess computers for our main measure of player performance (*Elo rating*), suggesting that AI in the form of chess computers helped chess players learn. Table 1 shows the effect of chess computers on player performance in the form of a binary treatment variable, *Chess computer access*. The average treatment effect on performance measured by the annual Elo rating is 0.011 with a p value smaller than .001 (Table 1, Model 1). Given that we use a linear model, the coefficient can be directly translated into Elo points. That is, the estimated effect represents an increase of about 11 Elo points in the player's annual Elo rating. This increase is economically significant, making a treated player with an initial Elo rating of 2000 about 3% more likely to beat an untreated player of previously equal rank under tournament conditions.

TABLE 1 Chess performance over time

Sample: All players	(1)	(2)	(3)	(4)	(5)
		Average centipawn loss			
DV	Elo rating	All	White	Black	Game won
Chess computer access	0.011 (0.003)	−0.362 (0.302)	−0.706 (0.342)	−0.381 (0.440)	0.060 (0.011)
Player FE	Yes	Yes	Yes	Yes	Yes
Player age FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Observations	170968	170968	151475	151475	93306
Players	19285	19285	17633	17633	5694
Log-likelihood	228875	−623057	−557466	−602732	−54990

Note: Columns (1) to (4) show the estimates of linear model regressions with high-dimensional fixed effects. The unit of observation is the individual chess player by year in columns (1) to (4). In column (5), the unit of observation is at the game level, where the sample consists of all games in which European and (former) Soviet players encountered each other. Note that a lower average centipawn loss means a higher chess performance. The samples in Columns (3) and (4) are aligned to ease comparison between the coefficients. Robust standard errors clustered at the player level in parentheses.

We test whether this effect is robust to our two alternative measures of performance.¹¹ We find the suggested negative effect on average centipawn loss, which implies that moves made by players with access to chess computers got closer to the best possible move available (see Table 1, Model 2). However, this effect is imprecisely estimated. We also find consistent evidence for our third dependent variable; that is, an increased likelihood that chess players with access to chess computers will beat those without it (see Table 1, Model 5).

To alleviate concerns that players from Western Europe and the Soviet Union are not comparable in skill development due to socio-economic differences between the regions, we run an event-study regression of player performance on a full set of interaction terms between the binary indicator for Western Europe and the year fixed effects as independent variables.

The event-study results presented in Figure 2 illustrate that the positive treatment effect on player performance is confined to the period when only Western players had access to chess computers. We find no systematic differences between Western and Soviet players before treatment (before 1977), a positive difference during the treatment period (1977–1989), and convergence back to a common trend afterwards (after 1989). This result solidifies the underlying linear trend assumption between the treatment and control groups during the two periods without treatment differences (before chess computers were available at all and after they became available in both regions).

¹¹We further test the robustness of these results by estimating the main specification on distinct subsamples. First, to ensure that our results do not hinge on the panel's unbalancedness and unequal density, we restrict the sample either to a more stable subset of players with above-median tournament activity (≥ 90 games) or to players with above-median career length (≥ 9 years) (see Supporting information Table B2, Models 1–4). Second, to address the concern that selection into our sample may drive the effect, we left-censor the panel to players with an Elo rating of 2000 or higher (see Table B2, Models 5 and 6). Third, we restrict the sample to shorter time windows around the two discrete events in 1977 and 1989 (see Table B3). These results all correspond qualitatively and statistically to our main result.




FIGURE 2 Event study: the effect of region (Western Europe) on chess performance (Elo rating). Source: This figure plots the point estimates and 95% confidence intervals of the effect of *Western Europe* $\times$ *Year* on Elo rating. The regression specification is as follows: $y_{it} = \alpha_i + \text{WesternEurope}_j \times \sum_{k=1970; k \neq 1973}^{2000} \beta_k + \text{age}_{it} + \delta_t$.

4.1.4 | Supporting evidence

While the staggered diffusion of chess computers provides the necessary exogenous variation in access, the treatment level remains crude (at the regional rather than player level). This leaves room for alternative explanations. We therefore conduct an additional test that strengthens the link between the observed increase in player performance and learning with chess computers. A particular limitation of early chess computers as artificial training partners was that they were more effective opponents when playing black (the second mover) instead of white (the first mover).¹² Consequently, players training with chess computers typically played white. In tournaments, however, a player is as likely to play black as white. If the observed improvement is indeed due to training with chess computers, one would expect players to have improved their white game more than their black game. To test this, we construct the average centipawn loss for each player's white and black tournament games separately. We indeed find that the effect on the average centipawn loss is larger for white games than for black games (Table 1, Models 3 and 4). Although the difference between these two coefficients has a p value of about .2, we deem it noticeable.¹³ This non-uniform effect rules out many non-

¹²Chess computers at that time lacked variation due to their deterministic play. Thus, when playing white, they would not vary their initial move as human players would do. This implies that chess computers constituted a better artificial training partner when the human player took the initiative.

¹³This is underlined by the results of an additional test (Supporting information Table B1), in which we regress chess computer access (as the dependent variable) on the two distinct average centipawn loss variables (as independent variables). In this regression, only the average centipawn loss variable based on white games shows a sizable negative coefficient with a p value of 0.05.




FIGURE 3 Chess computer strength over time. Source: This figure illustrates the progress in chess computer strength over time in Elo ratings. For each year, the upper bound strength of all available commercial chess computers is considered. Non-commercial chess computers (e.g., Deep Blue) are excluded.

technological explanations for the improved player performance in regions with access to chess computers.

Taken together, our results strongly suggest that AI-backed simulations in the form of chess computers help people learn how to engage in strategic interactions.

4.2 | Analysis—Part 2: AI strength as a boundary condition

We propose that decision-makers only benefit from training with AI if it provides artificial training partners whose skills match those of the respective trainee. This implies that a chess player benefits from training with chess computers only if they are (at least) on par with the chess player. To examine this, we leverage the continuous evolution of the strength of chess computers. Since their commercial introduction in 1977, chess computers have continuously improved in strength (see Figure 3). While early models played at best like an amateur, later models could also play at grandmaster level and above. If only players whose skills are inferior to current chess computers benefit from training with them, we expect to find heterogeneity in the treatment effect depending on whether current chess computers matched a player's skill.

We explore this heterogeneity in the treatment effect by including additional independent variables (and interactions thereof) that capture the strength of chess computers in absolute terms and in relative (i.e., player-specific) terms. *Chess computer strength* is the Elo rating of the strongest commercially available chess computer in a given year and *Chess computer superior* indicates whether the Elo rating of that chess computer is higher than that of the focal player for that year. The intuition behind *Chess computer strength* is that



TABLE 2 Chess performance over time—heterogeneity by chess computer strength

Sample: All players DV	Elo rating		Average centipawn loss	
	(1)	(2)	(3)	(4)
Chess computer (CC) access	0.015 (0.003)		0.050 (0.368)	
× CC strength	0.016 (0.005)		1.084 (0.601)	
× CC superior	0.038 (0.009)	0.040 (0.010)	0.074 (0.650)	0.217 (0.673)
× CC superior × CC strength	0.036 (0.009)	0.037 (0.009)	− 1.896 (0.766)	− 2.169 (0.771)
CC superior	− 0.053 (0.009)	− 0.055 (0.010)	− 0.328 (0.641)	− 0.503 (0.661)
Player FE	Yes	Yes	Yes	Yes
Player age FE	Yes	Yes	Yes	Yes
Year FE	Yes	No	Yes	No
Region-year FE	No	Yes	No	Yes
Observations	170968	170968	170968	170968
Players	19285	19285	19285	19285
Log-likelihood	229225	229309	− 623021	− 622907

Note: Columns (1) to (4) show the estimates of linear model regressions with high-dimensional fixed effects. The unit of observation is the individual chess player by year. Robust standard errors clustered at player level in parentheses.

trainees benefit more the stronger their artificial training partner is, while the intuition behind *Chess computer superior* is that chess computers superior to the player are particularly helpful as they can emulate human opponents that are at least on par with the player. The full model is:

$$\begin{aligned}
 y_{it} = & \alpha_i + \beta_1 \text{Chess computer access}_{it} + \beta_2 \text{Chess computer access} \times \text{Strength}_{it} \\
 & + \beta_3 \text{Chess computer superior}_{it} + \beta_4 \text{Chess computer access} \times \text{Superior}_{it} \\
 & + \beta_5 \text{Chess computer access} \times \text{Superior} \times \text{Strength}_{it} + \text{age}_{it} + \delta_t,
 \end{aligned} \quad (2)$$

where the dependent variable y is a player's performance, and the independent variables are *Chess computer access*, *Chess computer strength*, and *Chess computer superior* and their full set of interactions.¹⁴

We find that the magnitude of the treatment effect on player performance depends on the chess computer's strength in absolute terms as well as relative to the player (Table 2). In Model 1, the corresponding coefficients are positive, suggesting that chess computer access had a stronger effect on the performance of players to whom the best available chess computer was superior. Model 2 includes region-year fixed effects that capture time-variant differences between the two regions, which could be chess-specific (e.g., changes in a given chess player

¹⁴We further estimate the effect of chess computers on player performance with a more restrictive specification, in which we add region-calendar-year fixed effects, which capture differences in the year-specific chess environment between Western Europe and the Soviet Union. We can still estimate the effect of chess computer access because treatment now depends on the focal player's skill level and thus varies even between players within the same region and year.

population) or more general (e.g., changes in the socio-economic environment). The estimates of the two independent variables of interest are highly similar to those in the less-restrictive Model 1. In Models 3 and 4, we repeat this analysis with our other performance measure at the year-average-centipawn-loss level. The most relevant interaction estimates correspond to those in the models based on the Elo rating. That said, the coefficients of certain covariates (e.g., *Chess computer superior* and *Chess computer strength*) are estimated with much lower precision than in the previous models. In some cases, they even exhibit opposite signs.

These results suggest that a player benefits from access to chess computers predominantly if they can emulate a strategic response similar to that of a human opponent stronger than the focal player. This, in turn, implies that the worse the player, the more likely a chess computer of a given strength can be helpful.¹⁵ Note that this heterogeneity also helps us strengthen the identification of the main effect of chess computers, as we examine *within*-population variation.¹⁶

In summary, these findings of a heterogeneous treatment effect of chess computer strength on player performance are (a) methodologically helpful, as they corroborate the suggested effect of chess computers on player performance, and (b) theoretically interesting, as they imply a boundary condition—that is, the positive effect of AI on learning occurs only when the AI can provide sufficiently intelligent responses.

Our findings have a strategically important implication: not everybody benefits equally. Disadvantaged players—specifically, those with inferior skills—benefit most. AI-backed simulations may therefore reduce variance in skill distribution at the population level, with players in the left tail of the distribution more likely to improve. In other words, these disadvantaged players can catch up, resulting in a more level playing field.

We explore whether chess computers indeed changed the skill distribution by examining the distribution of Elo ratings of treated players over time. For this purpose, we focus on stable groups of Western and Soviet chess players: those active throughout the 1980s. While the variance of the skill distribution increases from 1980 to 1988 for both groups, we find—as expected—that this increase is *smaller* for Western players than for Soviet players (Supporting information Figure A8). This descriptive finding is consistent with the argument that chess computers fostered skill convergence and intensified competition.

4.3 | Analysis—Part 3: Mechanism

We propose that AI enables decision-makers to learn strategic interaction because artificial training partners resolve the bottleneck of scarce human training partners by constituting a scalable substitute. We therefore expect the effect of access to chess computers on player

¹⁵To show that the observed treatment effects are specific to players whose skills are inferior to those of the available chess computers in a given year, we run placebo regressions in which we manipulate the year-specific strength of available chess computers (see Figure A7). If we overstate the year-specific strength by 250 (500) Elo points, the coefficients of the independent variable *Chess computer access* $\times$ *Chess computer superior* $\times$ *Chess computer strength* become considerably smaller and can often no longer be statistically distinguished from zero. In contrast, if we choose a more conservative scenario and understate the year-specific strength of the available chess computers, the effects remain similar. These findings underscore the skill-level-dependent heterogeneity in treatment, as we would expect if only superior chess computers provide effective training.

¹⁶This heterogeneity also allows us to ascertain the effect of chess computers on players in both regions by splitting our sample and analyzing Western and Soviet players separately (see Supporting information Table B4).



performance to depend on the degree to which a given player has opportunities to train with human training partners (e.g., through friendly games in chess clubs). We therefore examine whether the availability of opportunities to train with humans moderates the effect of chess computer access on performance.

We keep performance as our dependent variable but introduce a moderator variable, *Chess event density*, which captures the availability of opportunities to train with humans, measured as the number of unique chess events in a player's country in the five previous years divided by the number of unique chess players over the same period. This variable is a time-variant country-specific measure of training opportunities and unrelated to chess computer access.

We indeed find that players with *more* training opportunities benefit *less* from access to chess computers (see Table 3). The interaction term in Model 1 (*Chess computer access* $\times$ *Chess event density*) suggests that a 1-SD-higher density of events in the player's home country lowers the baseline treatment effect by about 25%. In Models 2 and 3, we split our sample at the median into observations with a low and high chess event density and run separate regressions, as specified in Equation (2). The interaction coefficient (*Chess computer access* $\times$ *Chess computer superior* $\times$ *Chess computer strength*) is substantially larger for the low-density sample, suggesting that the effect of chess computer access on performance is stronger for players with constrained training opportunities. These results suggest that chess computers are a substitute for human training partners.

Our analysis suggests that a key mechanism by which AI helps people learn strategic interactions is providing training opportunities that would otherwise require (scarce) human training partners. Having leveraged the varying availability of opportunities to train with human opponents, we conduct a further test to see whether providing training is indeed the mechanism at work; that is, whether players with access to chess computers forgo training opportunities with humans. We examine whether they participate in fewer *Tournaments* and indeed find a corresponding effect (Table 3, Model 4), suggesting that chess players have substituted chess computer training for training at tournaments.

4.4 | Analysis—Part 4: imperfect substitution

Given the evidence that artificial training partners are a substitute for human training partners in learning strategic interaction, we examine whether they are a *perfect* substitute; that is, whether the learning outcomes of training with AI correspond to those of training with human training partners. We presume this not to be the case due to a distinct feature of AI—its consistency of performance. While human training partners frequently make idiosyncratic mistakes—so-called blunders—due to lack of concentration, limited memory, or varying mood, AI does not. Decision-makers who train with AI are therefore *not* exposed to as many blunders. We would thus expect players who train with AI to be worse at exploiting blunders than those without.

We indeed find that players with access to chess computers are less successful in exploiting an opponent's blunder to win the game. To analyze this, we return to the game level, choosing as our performance measure the incidence that a Western (Soviet) player wins in a direct encounter with a Soviet (Western) player (Table 4). The game-level analysis allows to establish a direct link between a blunder by the opponent and the outcome of the game. As a reference, we repeat the analysis from Section 4.2, where we find a positive effect of access to chess computers on winning (Model 1). We further find that blunders made by the

TABLE 3 Chess performance over time—mechanism: training opportunities

DV	(1) Elo rating	(2) Elo rating	(3) Chess event density	(4) Tournament attendance
Sample	All	Low	High	All
Chess computer (CC) access	0.012 (0.003)			
× Event density	−0.003 (0.001)			
× CC superior		0.040 (0.010)	0.040 (0.010)	−3.110 (0.519)
× CC superior × CC strength		0.041 (0.012)	0.027 (0.011)	−8.323 (0.866)
Event density	0.000 (0.001)			
CC superior		−0.055 (0.010)	−0.056 (0.010)	1.604 (0.502)
Player FE	Yes	Yes	Yes	Yes
Player age FE	Yes	Yes	Yes	Yes
Year FE	Yes	No	No	No
Region-year FE	No	Yes	Yes	Yes
Observations	170968	101969	97835	170968
Players	19285	16072	12832	19285
Log-likelihood	228900	136586	140080	−612445

Note: Columns (1) to (6) show the estimates of linear model regressions with high-dimensional fixed effects. Note that columns (2) and (3) are based on a sample split at the median of the chess event density for Western players, while Soviet players are part of both samples. The unit of observation is the individual chess player by year. Robust standard errors clustered at the player level in parentheses.

TABLE 4 Chess performance in case of blunder

Sample: All players	Game won		
DV	(1)	(2)	(3)
Chess computer (CC) access	0.060 (0.014)	0.057 (0.011)	0.077 (0.013)
Blunder by opponent		0.299 (0.006)	0.324 (0.013)
CC access × Blunder by opponent			−0.028 (0.013)
Player FE	Yes	Yes	Yes
Player age FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Observations	93306	93306	93306
Players	5694	5694	5694
Log-likelihood	−54990	−50874	−50871

Note: Columns (1) to (3) show the estimates of linear model regressions with high-dimensional fixed effects. The unit of observation is at game level, where the sample consists of all games in which European and (former) Soviet players encountered each other. Robust standard errors clustered at player level in parentheses.

opponent, which we operationalize as moves causing a centipawn loss in the 90th percentile, have a positive effect on winning (Model 2). However, once we add both independent variables—*Chess computer access* and *Blunders*—and their interaction (Model 3), we find



that players with access to chess computers benefit less from their opponents' blunders than players without access.¹⁷

This suggests that training with highly rational artificial training partners does not prepare trainees to exploit the blunders of boundedly rational human opponents as well as training with other humans does. Given that an important part of strategic interactions is to exploit opponents' blunders, AI-backed simulation remains limited in its ability to train decision-makers.

5 | DISCUSSION

Our analysis of *whether* AI helps decision-makers learn produces compelling evidence that access to chess computers improves players' performance. In particular, players with inferior skills and constrained training opportunities benefit from and are able to catch up due to training with AI. Our analysis reveals the underlying mechanism: AI-backed simulations provide artificial training partners that substitute for scarce human training partners and thus resolve the bottleneck created by that scarcity. Our comparison of the learning outcomes reveals, however, that artificial training partners are *imperfect* substitutes for human training partners. Because artificial training partners do not blunder, trainees do not learn to identify and exploit the blunders of the boundedly rational human opponents with whom they are training to compete.

5.1 | Organizational learning: AI as substitute for human training partners

Our research contributes to a key question in the learning literature: *What underlies the heterogeneity in learning progress?* Much research has illustrated the importance of accumulating experience (Anand et al., 2016; Argote et al., 2021; KC et al., 2013). But this raises the question of why actors differ in their tendency to accumulate experience. We illustrate that the scarcity of qualified human training partners is often a bottleneck in gaining experience when learning strategic interactions.

Our finding on training partners as a bottleneck and AI as a potential substitute informs research on the crucial role of social networks in learning. Actors learn from competitors (Posen et al., 2023; Ross & Sharapov, 2015), peers (Hallen et al., 2020; Tortoriello et al., 2012), teammates (Kim et al., 2023), partners (Ahuja, 2000; Powell et al., 1996), and others. An actor's network may thus be a source of competitive advantage or a bottleneck (Argote & Ingram, 2000). Our study illustrates that AI has the potential to be a substitute for network-based learning. It may thus undermine the role of networks as a source of competitive advantage.

Our research also informs the scalability of training interventions. Those that work at small scale often cannot be scaled because they rely on a non-scalable input (Levinthal & Wu, 2010; List, 2022). As Figure 1 illustrates, the particular role of AI-backed simulations is that they offer human-like responsiveness but are scalable. They may thus “democratize” responsive training.

Our finding that the learning outcomes of training with artificial and with human partners differ concerning the trainees' ability to recognize and exploit blunders has two important implications. First, while consistency in decision-making is generally considered a strength of AI,

¹⁷This result is robust to different operationalizations of blunders and to considering the actual number of blunders per game.

our findings suggest that, when training people, it can be a weakness, as it renders artificial training partners different from the boundedly rational humans with whom the trainee must eventually compete. Second, the consistency of AI decision-making implies—oddly enough—that AI-backed simulations programmed to make occasional human-like blunders may be better than their infallible counterparts at training decision-makers to compete against human competitors.

5.2 | AI in management and strategy: Training decision-makers

A key question for researchers and practitioners is how AI will affect various aspects of management and strategy (e.g., innovation (Cockburn et al., 2019; Wu et al., 2019), selection of ideas (Bell et al., 2023; Christensen et al., 2017), value creation and value capture (Iansiti & Lakhani, 2020), the organization of labor (Kane et al., 2021; Kellogg et al., 2020), decision-making (Glaeser et al., 2021; Lebovitz et al., 2022), or the provision of feedback (Tong et al., 2021)). In general research suggests that AI can *displace* or *complement* humans in task fulfillment (Brynjolfsson & Mitchell, 2017; Tschang & Almirall, 2021). Our study suggests that it can also help decision-makers *learn*—a so-far mostly neglected implication (Mollick & Mollick, 2022).

The insight that technologies can help decision-makers learn is not new (Edery & Mollick, 2008; Iyengar et al., 2015). What is new about training with AI is that it allows for the simulation of strategic interactions in which decision-makers receive human-like responses to their actions. It is thus plausible that, in the future, managers and employees will use AI to train for strategic (and perhaps other social) interactions.

We also provide an alternative view of how AI may help managers. Recent research shows the importance for managers of developing theory when navigating markets (Camuffo et al., 2021; Felin & Zenger, 2017; Gavetti et al., 2005; Zellweger & Zenger, 2022). Conceptual research shows AI's role in developing theory (Shrestha et al., 2021). We show that AI may also help decision-makers better understand strategic interaction. For example, rather than analyzing lots of data and revealing patterns to decision-makers, the AI serves as an artificial training partner allowing them to test their own strategies.

Our study also points to an alternative unit of analysis. While work on the effect of AI has focused on tasks, we focus on the decision-maker (also see Allen & Choudhury, 2022; Kang & Kim, 2023), finding that AI affects not only the fulfillment of tasks but also the people who fulfill them. As researchers assess the effects of AI on employee performance (Jia et al., 2023), it will be crucial to determine whether collaborating with AI has an effect on the individual that goes beyond the task at hand. It may, for example, increase actors' ability to reason and strategic intelligence (Helfat & Peteraf, 2015; Kang & Kim, 2023; Levine et al., 2017).

We illustrate that the relevance of AI for training depends on its state of evolution and on the trainee's skill level. Our results from chess computers suggest that *disadvantaged* actors may benefit first from AI in training and it seems plausible that the same pattern may occur in contexts in which the goal is not training but substitution. For example, AI-backed tools such as ChatGPT may initially be used to write college essays, but not (as yet) analytical journalism. However, the evolution of chess computers also shows the tremendous progress of AI to the point that it can be used to train even the very best.

We contribute to the emerging stream of *empirical* research on AI with a large-scale field-level historical dataset. While there has been rich debate on the implications of AI's



current state and its progress (Agrawal et al., 2019; Brynjolfsson & Mitchell, 2017; Felten et al., 2021), there has been little field data illustrating that progress. Chess is a domain for which such data is available. In addition to the long time horizon that allows measuring AI's increasing strength across decades, a particular feature is that the AI in chess can be compared with a human benchmark. Research on technology and strategy has illustrated the value of documenting how a technology's performance evolves (Adner et al., 2020; Adner & Kapoor, 2016; Casadesus-Masanell & Zhu, 2013), but such documentation has been missing for AI.

5.3 | Competitive strategy: Origin and dynamics of competitive advantage

Our analysis informs research on the origin and dynamics of competitive advantage. We complement research on the immense skill heterogeneity within industries (Kapoor & Agarwal, 2017; Lenox et al., 2006, 2010). Our results suggest that AI-backed simulation may foster convergence in the distribution of those skills learned best through responsive training. We found systematic effect heterogeneity indicating that *disadvantaged* players—specifically those with inferior skills or constrained training opportunities—benefitted most from chess computers. If disadvantaged actors benefit earlier and more, AI-based training can lead to a more equal skill distribution—leveling the playing field.

Our findings inform research on competitive moves. Research shows the importance of moves in strategic interactions (Chen et al., 2023; Dattée et al., 2018), but has focused on the consequences of competitive actions, leaving their antecedents rather unexplored. We show how managers may learn competitive moves.

Finally, we contribute to research on strategic interaction (Chen et al., 2010; Hannah & Eisenhardt, 2018; Van den Steen, 2018), illustrating the importance of decision-makers' ability to exploit the blunders of boundedly rational opponents.

5.4 | Generalizability, boundary conditions, and use cases

AI-backed simulations are not (yet) widespread, nor is it common (yet) for managers to use them for training strategic interactions. This lack of representativeness results, in part, from our choice of a context in which AI-backed simulation arrived *early*. While we do not suggest that our findings from this setting are representative (today), we do think that our study shows the potential of AI-backed simulation to train decision-makers and may thus be *indicative* of what will happen—or is already beginning to happen—in other settings.

It is very plausible that AI-backed simulation can help train decision-makers in settings comparable to chess. It is already used in training for other games, such as poker and Go (Choi et al., 2022; Kang et al., 2022; Riley, 2017; Shin et al., 2023). Non-game closed settings include multi-stage auctions (e.g., for mobile bandwidth licenses), in which firms may use AI-backed simulation to test-run bidding strategies.

Even in settings entirely unlike chess, it is plausible that decision-makers could use AI-backed simulation to train for strategic interactions. In such cases, the simulation is unlikely to correspond to the actual interaction. However, even if the simulation is only weakly correlated with the actual strategic interaction and its environment, the decision-maker is still likely to

benefit. It is also likely that the capabilities of AI-backed simulations will continue to develop so that they can be used to train strategic (and other social) interactions in ways that seem unlikely or even impossible today. For example, recent advances have made AI very powerful in negotiations (Kramár et al., 2022), which may allow human negotiators to use it to hone their skills. The emergence of large language models, like ChatGPT, makes it conceivable that AI can help people train in conversational skills (Mollick & Mollick, 2022). In fact, AI's increasing ability to master the Turing test indicates that artificial training partners could substitute for human training partners without the trainee even noticing.

More broadly, we expect economic and technological criteria to be decisive for widespread adoption of AI-backed simulation as a training mode. Concerning economic criteria, we expect the use of AI-backed simulation to be more likely where (a) training creates a high return, (b) experiential learning through actual interactions is costly, and (c) human training partners are particularly scarce. In the military, where strategic interactions are crucial but learning through actual interactions can cost lives, AI-backed simulations have already been tested.¹⁸

With respect to technological criteria, we expect that adopting AI-backed simulation requires AI sufficiently strong and free of algorithmic biases. First, as we show empirically, the training effect is conditional on the AI being strong enough to emulate competitors the trainee is likely to encounter in real life. Second, if the underlying algorithms are based on flawed data and/or code, the AI may be prone to systemic decision-making biases (Choudhury et al., 2020; Cowgill, 2019; Lambrecht & Tucker, 2019). These biases may be passed on to trainees, making them vulnerable to competitors who exploit the AI's weaknesses.

The requirement that the AI be sufficiently strong and unbiased has another implication: AI of that quality will in theory be able to replace the trainee. We therefore expect training with AI to be most relevant where human involvement remains socially desired (Glikson & Woolley, 2020), legally required (Tambe et al., 2019), or simply unavoidable (Brynjolfsson & McAfee, 2012). For example, even if AI could provide better diagnoses than a doctor can, people might still prefer to be diagnosed by a doctor.

5.5 | Limitations

Our research is subject to limitations. First, we assume homogeneity of treatment among chess players with access to chess computers, although players will have varying propensities to use them. We do so for lack of more-fine-grained data; we only observe whether chess players lived in a region where they *could* access chess computers, not whether they actually *did*. Despite abundant anecdotal evidence of chess computer use, we do not know which players used what kinds of chess computer and to what extent.¹⁹ This limitation most likely results in

¹⁸For instance, recent advances in AI have made combat simulators serious sparring partners. As one military instructor states, "In terms of emulating human reasoning, I feel this is to unmanned aerial vehicles what the IBM/Deep Blue versus Kasparov was to chess" (University of Cincinnati [2016], "New artificial intelligence beats tactical experts in combat simulation," available at <https://www.eurekalert.org/news-releases/782522>).

¹⁹Supporting information Figure A9 lists the main findings of a chess computer market survey of West German players in 1987. Among the estimated 9 million (amateur and professional) players, about 50% saw learning/training as the main reason to buy a chess computer. Moreover, the computer's performance was the most frequently stated primary criterion for purchase.



underestimation of the actual effect. Second, we rely on an admittedly coarse measure of training opportunities. As we have no data on a player's access to human training partners, we use the domestic density of chess events as a proxy. Third, we observe only the effect on the performance of professional and semiprofessional tournament players, not of amateurs. We would, however, expect—in line with our results—that amateurs benefited even more from chess computers, as the computers exceeded those players' skills early on.

5.6 | Managerial implications

Our research has managerial implications. Most notably, decision-makers may look for ways to use AI-backed simulation to train themselves or their employees. Our study also informs managers about those for whom such training is most likely helpful: relatively low-skilled employees and those for whom expensive training—such as training with human training partners—is not an option. Our study also provides managers with helpful insights about the limits of AI-backed simulations; they may not prepare people to spot and exploit real-life human blunders.

6 | CONCLUSION

Using the case study of chess computers, we examine the linkage between AI-backed simulation and learning strategic interactions. It is plausible that AI will democratize opportunities to learn strategic—and potentially more broadly social—interactions. If so, AI may become a game changer well beyond the game of chess. While it has often been asked whether AI will replace humans, our study suggests that a more imminent change may be that humans with AI will replace humans without AI.

ACKNOWLEDGMENTS

We are grateful for the help, suggestions and feedback we received from the editor and the reviewers as well as from Rajshree Agarwal, Vikas Aggarwal, Shahzad Ansari, Michaël Bikard, Linus Dahlander, Chuck Eesley, Kathleen Eisenhardt, Teppo Felin, Douglas Hannah, Dietmar Harhoff, Riitta Katila, Pooyan Khashabi, Xavier Martin, Anoop Menon, Daniel Levinthal, Juergen Mihm, Hart Posen, Alexey Rusakov, Evan Starr, Chris Tucci, Bala Vissa, Stefan Wagner, Phebo Wibbens, and Brian Wu, as well as from seminar participants at the Brigham Young University—Strategy Conference, Center for Advanced Study in the Behavioral Sciences, Columbia Business School, Copenhagen Business School, Cornell University, Dartmouth College, Duke University, EPFL, ESADE, ESMT Berlin, Harvard Business School, INSEAD, Pompeu Fabra University, Purdue University, St. Gallen University, Stanford University, Universidad Carlos III de Madrid, and the WhartonTech Conference.

DATA AVAILABILITY STATEMENT

The data that support the findings of this study are available via Chessbase (license), the World Chess Federation, and the website www.schach-computer.info.

ORCID

Fabian Gaessler  <https://orcid.org/0000-0001-6178-4097>

Henning Piezunka  <https://orcid.org/0000-0003-3956-1828>

REFERENCES

- Adner, R., Chen, J., & Zhu, F. (2020). Frenemies in platform markets: Heterogeneous profit foci as drivers of compatibility decisions. *Management Science*, 66(6), 2432–2451.
- Adner, R., Csaszar, F. A., & Zemsky, P. B. (2014). Positioning on a multiattribute landscape. *Management Science*, 60(11), 2794–2815.
- Adner, R., & Kapoor, R. (2016). Innovation ecosystems and the pace of substitution: Re-examining technology S-curves. *Strategic Management Journal*, 37(4), 625–648.
- Agrawal, A., Gans, J., & Goldfarb, A. (2018). *Prediction machines: The simple economics of artificial intelligence*. Harvard Business Press.
- Agrawal, A., Gans, J. S., & Goldfarb, A. (2019). Exploring the impact of artificial intelligence: Prediction versus judgment. *Information Economics and Policy*, 47, 1–6.
- Ahuja, G. (2000). Collaboration networks, structural holes, and innovation: A longitudinal study. *Administrative Science Quarterly*, 45(3), 425–455.
- Alcácer, J., Dezsó, C., & Zhao, M. (2015). Location choices under strategic interactions. *Strategic Management Journal*, 36(2), 197–215.
- Allen, R., & Choudhury, P. (2022). Algorithm-augmented work and domain experience: The countervailing forces of ability and aversion. *Organization Science*, 33(1), 149–169.
- Anand, J., Mulotte, L., & Ren, C. R. (2016). Does experience imply learning? *Strategic Management Journal*, 37(7), 1395–1412.
- Argote, L., & Ingram, P. (2000). Knowledge transfer: A basis for competitive advantage in firms. *Organizational Behavior and Human Decision Processes*, 82(1), 150–169.
- Argote, L., Lee, S., & Park, J. (2021). Organizational learning processes and outcomes: Major findings and future research directions. *Management Science*, 67(9), 5399–5429.
- Athey, S. C., Bryan, K. A., & Gans, J. S. (2020). The allocation of decision authority to human and artificial intelligence. *AEA Papers and Proceedings*, 110, 80–84.
- Bailey, D., Faraj, S., Hinds, P., von Krogh, G., & Leonardi, P. (2019). Special issue of organization science: Emerging technologies and organizing. *Organization Science*, 30(3), 642–646.
- Beane, M. (2019). Shadow learning: Building robotic surgical skill when approved means fail. *Administrative Science Quarterly*, 64(1), 87–123.
- Bell, J. J., Pescher, C., Tellis, G. J., & Fuller, J. (2023). Can AI help in ideation? A theory-based model for idea screening in crowdsourcing contests. *Marketing Science, Article in Advance*.
- Bloom, B. S. (1984). The 2 sigma problem: The search for methods of group instruction as effective as one-to-one tutoring. *Educational Researcher*, 13(6), 4–16.
- Brynjolfsson, E., Hui, X., & Liu, M. (2019). Does machine translation affect international trade? Evidence from a large digital platform. *Management Science*, 65(12), 5449–5460.
- Brynjolfsson, E., & McAfee, A. (2012). Winning the race with ever-smarter machines. *MIT Sloan Management Review*, 53(2), 53.
- Brynjolfsson, E., & Mitchell, T. (2017). What can machine learning do? Workforce implications. *Science*, 358(6370), 1530–1534.
- Camuffo, A., Gambardella, A., Messinese, D., Novelli, E., Paolucci, E., & Spina, C. (2021). A scientific approach to innovation management: Evidence from four field experiments. Working Paper.
- Casadesus-Masanell, R., & Zhu, F. (2013). Business model innovation and competitive imitation: The case of sponsor-based business models. *Strategic Management Journal*, 34(4), 464–482.
- Chase, W. G., & Simon, H. A. (1973). Perception in chess. *Cognitive Psychology*, 4(1), 55–81.
- Chen, E. L., Katila, R., McDonald, R., & Eisenhardt, K. M. (2010). Life in the fast lane: Origins of competitive interaction in new vs. established markets. *Strategic Management Journal*, 31(13), 1527–1547.
- Chen, T. X., Chen, J. H., Miller, D., Le Breton-Miller, I., & Chen, M. J. (2023). Corporate venture capital and interfirm rivalry: A competitive dynamics perspective. *Strategic Entrepreneurship Journal*, 17(1), 210–236.



- Choi, S., Kim, N., Kim, J., & Kang, H. (2022). How does AI improve human decision-making? Evidence from the AI-powered Go program. Working Paper.
- Choudhury, P., Starr, E., & Agarwal, R. (2020). Machine learning and human capital complementarities: Experimental evidence on bias mitigation. *Strategic Management Journal*, 41(8), 1381–1411.
- Christensen, K., Nørskov, S., Frederiksen, L., & Scholderer, J. (2017). In search of new product ideas: Identifying ideas in online communities by machine learning and text mining. *Creativity and Innovation Management*, 26(1), 17–30.
- Christensen, M., & Knudsen, T. (2010). Design of decision-making organizations. *Management Science*, 56(1), 71–89.
- Clough, D. R., & Piezunka, H. (2020). Tie dissolution in market networks: A theory of vicarious performance feedback. *Administrative Science Quarterly*, 65(4), 972–1017.
- Cockburn, I. M., Henderson, R., & Stern, S. (2019). The impact of artificial intelligence on innovation: An exploratory analysis. In A. Agrawal, J. Gans, & A. Goldfarb (Eds.), *The economics of artificial intelligence: An agenda* (pp. 115–146). University of Chicago Press.
- Courtney, H., Koller, T., & Lovallo, D. (2020). *Bias busters: War games? Here's what they're good for*. McKinsey Quarterly.
- Cowgill, B. (2019). Bias and productivity in humans and machines. Working Paper. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3584916
- Dattée, B., Alexy, O., & Autio, E. (2018). Maneuvering in poor visibility: How firms play the ecosystem game when uncertainty is high. *Academy of Management Journal*, 61(2), 466–498.
- Dell'Acqua, F. (2022). Falling asleep at the wheel: Human/AI collaboration in a field experiment on HR recruiters. Working Paper.
- Donskoy, M. V., & Schaeffer, J. (1988). Report on the 1st soviet computer-chess championship or re-awakening a sleeping giant. *ICGA Journal*, 11(2–3), 111–116.
- Edery, D., & Mollick, E. (2008). *Changing the game: How video games are transforming the future of business*. FT Press.
- Fang, C., & Levinthal, D. (2009). Near-term liability of exploitation: Exploration and exploitation in multistage problems. *Organization Science*, 20(3), 538–551.
- Felin, T., & Zenger, T. R. (2017). The theory-based view: Economic actors as theorists. *Strategy Science*, 2(4), 258–271.
- Felten, E., Raj, M., & Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, 42(12), 2195–2217.
- Garavetti, G., & Levinthal, D. (2000). Looking forward and looking backward: Cognitive and experiential search. *Administrative Science Quarterly*, 45(1), 113–137.
- Garavetti, G., Levinthal, D. A., & Rivkin, J. W. (2005). Strategy making in novel and complex worlds: The power of analogy. *Strategic Management Journal*, 26(8), 691–712.
- Glaeser, E. L., Hillis, A., Kim, H., Kominers, S. D., & Luca, M. (2021). Decision authority and the returns to algorithms. Working Paper.
- Glikson, E., & Woolley, A. W. (2020). Human trust in artificial intelligence: Review of empirical research. *Academy of Management Annals*, 14(2), 627–660.
- Greve, H. R. (2003). *Organizational learning from performance feedback: A behavioral perspective on innovation and change*. Cambridge University Press.
- Gustafson, T. (1981). *Selling the Russians the Rope? Soviet technology policy and US export control*. Rand Corporation.
- Hallen, B. L., Cohen, S. L., & Bingham, C. B. (2020). Do accelerators work? If so, how? *Organization Science*, 31(2), 378–414.
- Hannah, D. P., & Eisenhardt, K. M. (2018). How firms navigate cooperation and competition in nascent ecosystems. *Strategic Management Journal*, 39(12), 3163–3192.
- Helfat, C. E., & Peteraf, M. A. (2015). Managerial cognitive capabilities and the microfoundations of dynamic capabilities. *Strategic Management Journal*, 36(6), 831–850.
- Iansiti, M., & Lakhani, K. R. (2020). *Competing in the age of AI: Strategy and leadership when algorithms and networks run the world*. Harvard Business Press.

- Iyengar, K., Sweeney, J. R., & Montealegre, R. (2015). Information technology use as a learning mechanism. *MIS Quarterly*, 39(3), 615–642.
- Jia, N., Luo, X., Fang, Z., & Liao, C. (2023). When and how artificial intelligence augments employee creativity. *Academy of Management Journal*, *In Press*.
- Joseph, J., Klingebiel, R., & Wilson, A. J. (2016). Organizational structure and performance feedback: Centralization, aspirations, and termination decisions. *Organization Science*, 27(5), 1065–1083.
- Judy, R. W., & Clough, R. W. (1990). Soviet computing in the 1980s: a survey of the software and its applications. *Advances in Computers*, 30, 223–306.
- Kane, G. C., Young, A. G., Majchrzak, A., & Ransbotham, S. (2021). Avoiding an oppressive future of machine learning: A design theory for emancipatory assistants. *MIS Quarterly*, 45(1), 371–396.
- Kang, J., Yoon, J. S., & Lee, B. (2022). How AI-based training affected the performance of professional Go players. Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems, 1–12.
- Kang, X., & Kim, H. (2023). Predictive algorithms and decision-making: Evidence from a field experiment. Working Paper.
- Kapoor, R., & Agarwal, S. (2017). Sustaining superior performance in business ecosystems: Evidence from application software developers in the iOS and Android smartphone ecosystems. *Organization Science*, 28(3), 531–551.
- Kasparov, G. (2018). Chess, a drosophila of reasoning. *Science*, 362(6419), 1087.
- KC, D., Staats, B. R., & Gino, F. (2013). Learning from my success and from others' failure: Evidence from minimally invasive cardiac surgery. *Management Science*, 59(11), 2435–2449.
- Keil, T., Posen, H., & Workiewicz, M. (2022). Aspirations, beliefs, and a new idea: Building on March's other model of performance feedback. *Academy of Management Review*, *In Press*.
- Kellogg, K. C., Valentine, M. A., & Christin, A. (2020). Algorithms at work: The new contested terrain of control. *Academy of Management Annals*, 14(1), 366–410.
- Kim, S. H., Song, H., & Valentine, M. A. (2023). Learning in temporary teams: The varying effects of partner exposure by team member role. *Organization Science*, 34(1), 433–455.
- Krakowski, S., Luger, J., & Raisch, S. (2022). Artificial intelligence and the changing sources of competitive advantage. *Strategic Management Journal*, 44(6), 1425–1452.
- Kramár, J., Eccles, T., Gemp, I., Tacchetti, A., McKee, K. R., Malinowski, M., Graepel, T., & Bachrach, Y. (2022). Negotiation and honesty in artificial intelligence methods for the board game of Diplomacy. *Nature Communications*, 13(1), 7214.
- Lambrecht, A., & Tucker, C. (2019). Algorithmic bias? An empirical study of apparent gender-based discrimination in the display of STEM career ads. *Management Science*, 65(7), 2966–2981.
- Lebovitz, S., Lifshitz-Assaf, H., & Levina, N. (2022). To engage or not to engage with AI for critical judgments: How professionals deal with opacity when using AI for medical diagnosis. *Organization Science*, 33(1), 126–148.
- Lenox, M. J., Rockart, S. F., & Lewin, A. Y. (2006). Interdependency, competition, and the distribution of firm and industry profits. *Management Science*, 52(5), 757–772.
- Lenox, M. J., Rockart, S. F., & Lewin, A. Y. (2010). Does interdependency affect firm and industry profitability? An empirical test. *Strategic Management Journal*, 31(2), 121–139.
- Levine, S. S., Bernard, M., & Nagel, R. (2017). Strategic intelligence: The cognitive capability to anticipate competitor behavior. *Strategic Management Journal*, 38(12), 2390–2423.
- Levinthal, D. A., & Wu, B. (2010). Opportunity costs and non-scale free capabilities: Profit maximization, corporate scope, and profit margins. *Strategic Management Journal*, 31(7), 780–801.
- List, J. A. (2022). *The voltage effect: How to make good ideas great and great ideas scale*. Crown Publishing Group.
- Marino, A., Aversa, P., Mesquita, L., & Anand, J. (2015). Driving performance via exploration in changing environments: Evidence from Formula One racing. *Organization Science*, 26(4), 1079–1100.
- Martignoni, D., Keil, T., & Lang, M. (2020). Focus in searching core-periphery structures. *Organization Science*, 31(2), 266–286.
- Martin, B. O., Kolomitro, K., & Lam, T. C. (2014). Training methods: A review and analysis. *Human Resource Development Review*, 13(1), 11–35.



- Maula, M., Heimeriks, K. H., & Keil, T. (2022). Organizational experience and performance: A systematic review and contingency framework. *Academy of Management Annals*, *In Press*.
- McCorduck, P. (2004). *Machines who think: A personal inquiry into the history and prospects of artificial intelligence*. CRC Press.
- Menon, A. R., & Yao, D. A. (2017). Elevating repositioning costs: Strategy dynamics and competitive interactions. *Strategic Management Journal*, *38*(10), 1953–1963.
- Mollick, E. R., & Mollick, L. (2022). New modes of learning enabled by AI chatbots: Three methods and assignments.
- Panico, C. (2017). Strategic interaction in alliances. *Strategic Management Journal*, *38*(8), 1646–1667.
- Piezunka, H., Aggarwal, V. A., & Posen, H. E. (2022). The Aggregation–Learning Trade-off. *Organization Science*, *33*(3), 1094–1115.
- Posen, H. E., Ross, J. M., Wu, B., Benigni, S., & Cao, Z. (2023). Reconceptualizing imitation: Implications for dynamic capabilities, innovation, and competitive advantage. *Academy of Management Annals*, *17*(1), 74–112.
- Powell, W. W., Koput, K. W., & Smith-Doerr, L. (1996). Interorganizational collaboration and the locus of innovation: Networks of learning in biotechnology. *Administrative Science Quarterly*, *41*(1), 116–145.
- Rahmandad, H., Repenning, N., & Sterman, J. (2009). Effects of feedback delay on learning. *System Dynamics Review*, *25*(4), 309–338.
- Riley, T. (2017). Artificial intelligence goes deep to beat humans at poker. *Science*, *16*, 37–43.
- Rivkin, J. W., & Siggelkow, N. (2003). Balancing search and stability: Interdependencies among elements of organizational design. *Management Science*, *49*(3), 290–311.
- Ross, J. M., & Sharapov, D. (2015). When the leader follows: Avoiding dethronement through imitation. *Academy of Management Journal*, *58*(3), 658–679.
- Salas, E., Tannenbaum, S. I., Kraiger, K., & Smith-Jentsch, K. A. (2012). The science of training and development in organizations: What matters in practice. *Psychological Science in the Public Interest*, *13*(2), 74–101.
- Sharapov, D., & Ross, J. M. (2023). Whom should a leader imitate? Using rivalry-based imitation to manage strategic risk in changing environments. *Strategic Management Journal*, *44*(1), 311–342.
- Shin, M., Kim, J., van Opheusden, B., & Griffiths, T. L. (2023). Superhuman artificial intelligence can improve human decision-making by increasing novelty. *Proceedings of the National Academy of Sciences*, *120*(12), e2214840120.
- Shrestha, Y. R., He, V. F., Puranam, P., & von Krogh, G. (2021). Algorithm supported induction for building theory: How can we use prediction models to theorize? *Organization Science*, *32*(3), 856–880.
- Siggelkow, N. (2002). Evolution toward fit. *Administrative Science Quarterly*, *47*(1), 125–159.
- Simon, H. A. (1996). *The sciences of the artificial*. MIT Press.
- Simon, H. A., & Schaeffer, J. (1992). The game of chess. In R. J. Aumann & S. Hart (Eds.), *Handbook of game theory with economic applications* (pp. 1–17). Elsevier.
- Smith, K. G., Ferrier, W. J., & Ndofor, H. (2005). Competitive dynamics research: Critique and future directions. In M. A. Hitt, R. E. Freeman, & J. S. Harrison (Eds.), *The Blackwell handbook of strategic management* (pp. 309–354). Blackwell Publishing.
- Tambe, P., Cappelli, P., & Yakubovich, V. (2019). Artificial intelligence in human resources management: Challenges and a path forward. *California Management Review*, *61*(4), 15–42.
- Tong, S., Jia, N., Luo, X., & Fang, Z. (2021). The Janus face of artificial intelligence feedback: Deployment versus disclosure effects on employee performance. *Strategic Management Journal*, *42*(9), 1600–1631.
- Tortoriello, M., Reagans, R., & McEvily, B. (2012). Bridging the knowledge gap: The influence of strong ties, network cohesion, and network range on the transfer of knowledge between organizational units. *Organization Science*, *23*(4), 1024–1039.
- Tschang, F. T., & Almirall, E. (2021). Artificial intelligence as augmenting automation: Implications for employment. *Academy of Management Perspectives*, *35*(4), 642–659.
- van den Steen, E. (2018). The strategy in competitive interactions. *Strategy Science*, *3*(4), 574–591.
- Wu, L., & Brynjolfsson, E. (2015). The future of prediction: How Google searches foreshadow housing prices and sales. In A. Goldfarb, S. M. Greenstein, & C. E. Tucker (Eds.), *Economic analysis of the digital economy* (pp. 89–118). University of Chicago Press.

- Wu, L., Lou, B., & Hitt, L. (2019). Data analytics supports decentralized innovation. *Management Science*, 65(10), 4863–4877.
- Zellweger, T., & Zenger, T. (2022). Entrepreneurs as scientists: A pragmatist alternative to the creation-discovery debate. *Academy of Management Review*, 47(4), 696–699.

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

How to cite this article: Gaessler, F., & Piezunka, H. (2023). Training with AI: Evidence from chess computers. *Strategic Management Journal*, 44(11), 2724–2750.
<https://doi.org/10.1002/smj.3512>