TRAINING WITH AI -

COMPETING AGAINST HUMANS

Fabian Gaessler
Pompeu Fabra University

Henning Piezunka
INSEAD

Strategic interactions are difficult to learn due to the need and scarcity of training partners. We suggest that AI can help decision-makers learn strategic interactions by serving as *artificial* training partners. We present evidence from chess computers, the first widespread incarnation of AI. Leveraging their staggered diffusion, we find that chess computers indeed helped chess players improve their performance as a substitute for scarce human training partners. We also illustrate that chess computers were *not* a perfect substitute, as players training with them did not get exposed to idiosyncratic ("human") mistakes and thus did not learn to exploit them. We discuss implications for research on organizational learning, on AI in management and strategy, and on competitive advantage.

To succeed in competition, organizations and their decision-makers need to master strategic interactions (e.g., airlines setting prices and routes) (Schelling, 1960; Axelrod, 1984; Ferrier, Smith, and Grimm, 1999; Chen, Katila, McDonald, and Eisenhardt, 2009; Chen and Miller, 2012; van den Steen, 2018). Given the high stakes in strategic interactions, learning from actual experience is expensive. Decision-makers may thus choose to simulate the competition with the help of training partners who emulate the competitors (e.g., CEOs hire consultants to engage in "war games"). However, training partners that are qualified to do so are often scarce. The limited availability of training partners may thus constitute a bottleneck for decision-makers who strive to learn strategic interactions.

We suggest that AI can help decision-makers learn strategic interactions to compete against humans. Specifically, we suggest that AI elevates computer-based simulations. Within these simulations, the AI serves as an *artificial training partner* that emulates competitors by providing intelligent responses to the actions of the decision-maker. This allows decision-makers to learn as they can try out different actions and see how a competitor would likely respond. Computer-based simulations where the AI emulates the competitors thus serve as a scalable *substitute* to oftentimes scarce human training partners. In brief, we propose that AI helps decision-makers to overcome the scarcity of training partners, and thus allows them to learn strategic interactions.

We test our theory about the potential to learn strategic interaction with the help of AI by examining the effect of chess computers on the performance of chess players. This empirical setting allows us to link the first widespread application of AI (i.e., chess computers) with a prime example of strategic interaction (i.e., chess). We use data on more than 50,000 chess players Western Europe and the (former) Soviet Union and analyze their performance across 1.6 million tournament games. Our window of observation reaches from 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.

We start out by investigating the *main effect*; that is, whether chess computers have a positive effect on chess players' performance. In a DiD framework, we leverage two natural experiments that provide exogenous 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 (former) 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 players in the (former) Soviet Union from 1989. We conduct additional analyses that provide compelling evidence that the increase in performance is indeed due to chess computers. First, we exploit a particular limitation of early chess computers, which rendered them more effective opponents in games where the player is "white". If the positive effect on performance is due to training with chess computers, one would therefore expect that players have improved their "white" game more than their "black" game. We find this to be the case. Second, we leverage that chess computers have continuously improved in strength since their introduction in 1977. While early models played at best like an amateur, later models played on grandmaster level and above. We find that the effect on performance is confined to players whose skills are inferior to the chess computers available at a given time. This within-group treatment heterogeneity also suggests an important boundary condition: the AI needs to be sufficiently intelligent to provide effective training partners.

We then test the suggested *mechanism*, that is, whether artificial training partners constitute a substitute for human training partners and thus help trainees overcome the scarcity of human training partners. To distinguish players by how constrained their training opportunities with human training partners are, we leverage the density of local chess events. We find that players constrained in their training opportunities benefit *more* from chess computers. By extension, our argument further suggests that players with access to chess computers may forgo participation in tournaments that served primarily as an opportunity for them to train with other humans. Our data provides consistent evidence: treated players participate in fewer tournaments.

Despite the evidence that artificial training partners are a substitute for human training partners in learning strategic interaction, it is not clear whether they constitute a *perfect substitute*. We seek an answer to this by examining *whether the learning outcome differs between training with AI and training with a human*. We thereby focus on a distinct feature of AI: the consistency of their performance. While boundedly rational human training partners conduct idiosyncratic mistakes—so-called blunders—due to lack of concentration, the AI does not as “machines don’t get tired”. Decision-makers who train with AI thus have less exposure to blunders. Part of succeeding in competitive strategic interaction is, however, to identify and exploit blunders of one opponent. We indeed find that players with access to chess computers are less successful in exploiting an opponent’s blunder when competing against humans.

By examining how decision-makers learn strategic interactions, we contribute to research on *learning*. While a rich body of research has examined the learning of various kinds of activities (Darr, Argote, and Epple, 1995; Argote, 2012; KC, Staats, and Gino, 2013; Anand, Mulotte, and Ren, 2016; Maula, Heimeriks and Keil, 2022), there has been a paucity of research on the learning of strategic interactions. A key insight that emerges from our study of learning strategic interactions is that responsive training partners constitute a bottleneck for many decision-makers. This bottleneck may help explain the heterogeneity in actors’ tendency to train and to accumulate experience. Our findings also illustrate that it is possible to overcome this bottleneck as AI can substitute for human training partners. Whether and where AI will be a meaningful substitute for human training partners beyond chess and similar settings remains to be seen and depends in part on further progress in AI. However, if it does, it will have tremendous consequences as it renders responsive training of strategic interactions more accessible, which has been shown to provide highly effective learning (Bloom, 1984). Finally, our finding that and how learning with artificial instead of human training partners differs illustrates how learning in a testbed may ill-prepare decision-makers to act in the real case; competing against humans is different from competing against AI.

We also contribute to research on the role of AI in management and strategy. While research has pointed out that AI can substitute and complement human decision-makers (Brynjolfsson and Mitchell, 2017; Acemoglu and Restrepo, 2018; Agrawal, Gans, and Goldfarb, 2019), we show that it can also help them learn by resolving the trade-off between availability and responsiveness that defines the landscape of conventional training modes. It does so by combining the *scalability* of computer-based simulations and the ability to respond intelligently of human training partners. Our account is thus consistent with prior work as artificial training partners fulfill the same role as human training partners and thus pose a substitute for them. However, it also contributes by illustrating a so-far neglected implication: the role of AI in training people. We also illustrate that the relevance of AI for training purposes depends on the state of evolution of the AI *and* the skill-level of the trainee. Our results suggest that weaker actors are the ones who may benefit first from AI in training, but eventually that AI has the potential to train even the very best. An interesting feature is that the consistency of AI - typically seen as a strength - may also constitute a weakness when it trains humans in strategic interaction: they do not get exposed to idiosyncratic ("human") mistakes and thus do not learn to exploit them. We also make an empirical contribution, while conceptual research has made enormous progress in examining the future role of AI, we complement this debate by extensive historical data on an industry where AI has already been a game changer.

Finally, our findings inform research on the origin and dynamics of competitive advantage. First, we illustrate that access to human training partners can constitute a source of competitive advantage for decision-makers. Prior research that examines the role of people the focal actor has access to, has typically emphasized their role as collaborateurs, trading partners, or sources of information (Burt, 2005; Uzzi and Spiro, 2005); our emphasis on training partners points to a so far neglected function. Second, we complement prior research that has examined population-level distribution in performance (Lenox, Rockart, and Lewin, 2006, 2010; Porter and Siggelkow, 2008; Kapoor and Agarwal, 2017) by illustrating how new training technologies can systematically change it. Specifically, we find that the skill distribution among players with access to chess computers became more narrow and, consequently, the competition more intense as disadvantaged players (i.e., players of inferior skills or who lack training partners) benefitted more and caught up. We also contribute to research on competitive moves: while research has pointed to the importance of moves in strategic interactions (Ferrier, Smith, and Grimm, 1999; Smith, Ferrier, and Ndofor, 2005), our study illustrates how decision-makers may learn such competitive moves. Finally, we contribute to research on strategic interaction: we illustrate that decision-makers' ability to exploit blunders of their boundedly rational opponents constitute an important skill—a skill that is better acquired when training with boundedly humans than with consistently high-performing machines.

2 THEORY

2.1 The difficulty of learning strategic interactions

Engaging in strategic interaction with a competitor is different from other kinds of organizational actions—and so is the process of learning it. While typically taking an action is followed by some kind of performance feedback (Greve, 2003; Gaba and Joseph, 2013), in competitive strategic interactions, the focal actor engages with a competitor in a sequence of actions and *responses*—and only at their end stands the actual outcome (Schelling, 1960; Smith, Ferrier, Ndofor, 2005; Chen, Katila, McDonald, and Eisenhardt, 2009; Gavetti and Menon, 2016; Menon and Yao, 2017; Levine, Bernard, and Nagel, 2017).[1] To succeed in strategic interactions thus requires decision-makers to anticipate their counterparts' responses to their actions as the eventual outcome depends on them. The sequence of actions and responses (and the need to anticipate it) renders the process of learning strategic interactions particularly complicated.

Prior research on (organizational) learning points to specific complications that result from the sequence of actions and responses that precedes the eventual outcome. For example, the sequence of actions and responses—each involving choices—results in a large number of scenarios, making it hard to gain an overview of the performance landscape (Simon, 1947; March and Simon, 1958; Csaszar, 2018). Also, the effectiveness of actions is strongly contingent on other actions and responses, resulting in intricate interdependencies that add to the difficulty of learning (Levinthal, 1997; Siggelkow, 2002; Baumann and Siggelkow, 2013; Levine, Bernard, and Nagel 2017; Rahmandad, 2019). Moreover, credit assignment is difficult because performance is evaluated only at the end of the interaction (Denrell, Fang, and Levinthal, 2004; Fang and

[1] For example, prior literature has considered location choices and alliance building as strategic interactions at firm-level (Alcacer, Dezső, and Zhao, 2015; Panico, 2016). More generally, any bargaining between firms (e.g., auctions, litigation) or within firms (e.g., salary negotiations) can be understood as a strategic interaction.

Levinthal, 2009; Rahmandad, Repenning, and Sterman, 2009). 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 information (e.g., in books or lectures). They thus learn about potential actions and responses but do not go through the experience of taking and receiving them (Martin, Kolomitro, and Lam, 2014). While learning from codified information is helpful as learners gain a general understanding and familiarize themselves with archetypical strategies, it is unlikely to be sufficient to master strategic interactions. Even if the codification is extensive, in actual strategic interactions decision-makers unavoidably encounter non-codified scenarios. Already slight deviations from codified scenarios can result in fundamentally different outcomes that require different actions. In contrast, experiential learning can help actors to become accustomed to navigating unknown scenarios by improving their ability to abstract and develop heuristics. Learning from codified information is also insufficient insofar as succeeding in strategic interaction goes beyond understanding: it also requires implementation. The challenges in learning strategic interactions and the insufficiencies of learning them in an non-experiential manner suggests that experiential learning is essential.

A rich body of work illustrates the benefits of experiential learning (Darr, Argote, and Epple, 1995; Anand, Mulotte, Ren, 2016; Huckman, Staats, and Upton, 2009; Maula, Heimeriks, and Keil, 2022; Schilling, Vidal, Ployhart, and Marangoni, 2003), and these benefits apply (particularly) when learning strategic interactions. Learning strategic interaction experientially implies that the person takes an action and receives a response by an interaction partner. The response helps decision-makers to learn about the effectiveness of their actions; it leads to a better understanding of potential responses and the outcomes associated with an action.[2] Decision-makers can also explore alternative actions, gaining an understanding of the *relative* effectiveness and extending their repertoire of actions (Ferrier, Smith, and Grimm, 1999; Smith, Ferrier, and Ndofor, 2005). Taken together, experiential learning where decision-makers receive a response to their actions allows them to improve their performance in strategic interactions as they become better in developing and understanding a repertoire of actions.

2.4 Traditional ways of learning strategic interactions experientially

One way to learn experientially is *to engage in the actual strategic interaction*; i.e., to learn "online" (Gavetti and Levinthal, 2000). Actual engagement in strategic interaction may mean, for example, for airline executives to try out competitive actions like adjusting prices, launching new routes, etc., and then observe the corresponding responses by their competitor as well as the eventual outcome of the strategic interaction. While engaging in the actual strategic interaction allows decision-makers to learn from experience, it comes with several downsides: 1) opportunities to engage in actual strategic interaction are often limited; 2) the high stakes of actual strategic interaction can make exploration with uncertain outcome expensive; and 3) engaging in the actual strategic interaction may have undesired externalities as it educates competitors about the competitive landscape. In brief, engaging in actual strategic interactions is very costly.

---

[2] Experiential learning allows decision-makers to reduce the frequency of false positives and false negatives when choosing actions (Gavetti and Levinthal, 2000). With respect to false positives, the decision-maker tries out an action only to realize after receiving a signal that the action is not as effective as assumed. For example, consider airline executives who cut prices expecting to grab a bigger market share, but instead trigger a price war. With respect to false negatives, decision-makers observe the responses they receive and may realize that the response also constitutes a viable action for them. For example, consider manufacturing executives who invest in a marketing campaign to increase demand, only to realize that their competitor has contractually secured most of the supplier capacity, limiting the focal firm's ability to satisfy the increased demand. This mishap may nevertheless teach them the strategy of preempting a competitor upstream.

In light of the benefits of experiential learning as well as the costs of engaging in *actual* strategic interactions, decision-makers seek to learn strategic interactions experientially by *simulating the strategic interaction with human training partners*. For example, executives of a firm hire a consulting company to engage in war games where they explore actions—and consultants emulate potential (and ideally realistic) responses by the actual competitors. While simulating the strategic interactions with the help of human training partners who emulate the competition allows for experiential learning, it comes with various challenges and downsides. Human training partners are oftentimes scarce. This is because training with a human training partner is usually bilateral (Bloom, 1984); i.e., not scalable. Moreover, human training partners may be hard to find and/or costly given their opportunity costs (they may gain little from training others). This is particularly true for human training partners that are sufficiently skilled to emulate the response of actual competitors.[3] In brief, training partners are scarce, and thus constitute a potential bottleneck in learning strategic interaction.

2.3 Learning strategic interactions experientially with AI

We suggest that AI can serve as a *substitute* for training with training partners and thus help decision-makers to overcome the outlined bottlenecks and help them to learn strategic interactions. We conceptualize the AI to be part of a computer-based simulation[4]; the computer-based simulation provides the environment of the strategic interaction, and the AI emulates the (responses of the) competitor. The specific role of the AI in such simulations is thus to emulate a competitor and its responses—and thus serve as an *artificial training partner*.

[3] The importance and difficulty of finding training partners that are able to emulate competitors' responses becomes evident in the approach of the Chinese military: it hired retired British fighter jet pilots to emulate potential responses of the opponent in a conflict.
[4] Conventional computer-based simulations have long been used to train people in certain tasks thanks to their ability to simulate the environment. For example, pilots train with flight simulators (Leemkuil and de Jong, 2012), where the physical environment of the airplane can be realistically simulated even without AI.

AI helps to elevate computer-based simulations by emulating competitors in a more realistic manner than previously possible. While conventional simulations sometimes also involve competitors, they fall short to emulate them realistically (cf. Salas, Tannenbaum, Kraiger, and Smith-Jentsch, 2012). This is because when conventional simulations strive to emulate human responses, they follow non-complex rules and patterns, which rarely reflect actual human decision-making. Indeed, such simulations often give the interaction partners particular advantages (e.g., superior knowledge about the environment) to make them more challenging competitors despite their lack of intelligence. By contrast, AI-featuring computer-simulation can realistically emulate human behavior. Advances in computer algorithms, such as neural networks, combined with ever-increasing computational power and data availability have made it possible to teach computers to make intelligent decisions in a way that is similar to how humans would make them - that is, to create an artificial intelligence (Agrawal et al., 2018; Brynjolfsson and McAfee, 2014). Although the current stage of AI varies substantially by domain, we suggest that AI allows the depiction of human decision-making necessary to provide an intelligent response to a focal actor in strategic interactions.

An example provided by a military training expert illustrates the difference AI makes: Consider a simulation where soldiers are tasked to conquer a territory. Without AI, the various emulated enemies each act according to fairly simple and predefined routines. In other words, they do not realistically respond to the soldiers' actions. With AI, the emulated enemies coordinate among one another and respond to the soldiers' actions as real enemies would do. This difference can be crucial for learning as it separates a toy-like simulation from a sufficiently realistic training simulation (Grossman, Heyne, and Salas, 2015). In brief, AI makes it feasible that computer-based simulations emulate intelligent responses of humans—and can thus help to train decision-makers in strategic interactions.

The particular appeal of AI-featuring simulation is thus that 1) it offers experiential learning that is comparable to training with (responsive) human training partners, and 2) it is scalable and thus highly available (as conventional computer-based simulations). Figure 1 compares training with AI to other modes of training in terms of responsiveness and availability, illustrating that AI resolves the trade-off between responsiveness and availability that defines the landscape of conventional training modes. Beyond helping to position training with AI, the figure also illustrates why training with AI is not just beneficial when compared to no training, but also above and beyond other training modes.

[INSERT FIGURE 1 ABOUT HERE]

The arguments suggest various *implications*. First and foremost, we expect "a main effect" that decision-makers benefit from simulations with AI so that their performance in strategic interactions improves. It is noteworthy that the argument above does not simply suggest that training with AI is better than no training, but also that it can be beneficial "in the field", where other modes of training strategic interactions already exist.

Beyond this "main effect" that AI can help decision-makers improve, we also expect an interaction between training with AI-featuring simulations and other ways of learning strategic interactions experientially. Specifically, we have outlined the mechanism that AI helps decision-makers improve because scalable artificial training partners resolve the bottleneck of scarce human training partners. AI may thus have the largest effect where decision-makers have few opportunities to train with human training partners. Relatedly, we have argued that engaging in actual strategic interaction can be costly. AI may thus reduce the engagement of decision-makers in actual strategic interactions, which only served training purposes and can be substituted with training with artificial training partners.

In our line of reasoning, we build on the assumption that AI-featuring simulations serve as a substitute for scarce human-training partners. However, it is not clear to what extent this assumption is actually fulfilled. We therefore examine whether AI-featuring simulations constitute a "perfect substitute": that is whether what is learnt when training with AI is identical to what is learnt when training with human training partners.

3 SETTING AND DATA: CHESS COMPUTERS

In order to examine whether AI can help people improve their performance in strategic interactions, we study the effect of chess computers on the performance of (human) chess players. This empirical setting allows us to link the first widespread application of AI—chess computers—with a prime example of strategic interaction.

3.1 Chess as Strategic Interaction, Chess Computers as AI-featuring Simulations

We use chess to examine decision-makers' ability to learn strategic interactions. In chess, two actors encounter each other in a closed setting, in which each actor takes an action followed by the opponent's response. When players consider their next move, they try to anticipate their opponent's response in each scenario. In fact, a player's chess skill is typically equated to the ability to anticipate a long chain of actions that follows from a given move. Given the interactive nature of the game where players take actions and respond to one another, the large number of scenarios, and the difficulty of credit assignment, chess is widely considered a prototypical—and often-cited—example of strategic interaction (Wernerfelt, 1995; Powell, 2003; Puddephatt, 2003; Krakowski, Luger, and Raisch, 2022), which stands out due to its closed setting and high frequency of interactions.

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, father of modern computer science, and Claude Shannon, father of modern information theory, both turned their attention to chess (Shannon, 1950; Turing, 1953). For decades, chess has been the primary use case for AI innovation and chess computers have been at the frontier of AI applications. Several milestones in AI history—first program, first human-machine competition, first superintelligence, first reinforced learning algorithm—involved chess (Ensmenger, 2011; Kasparov, 2018). Herbert Simon, a pioneer of AI (and of organization theory), described chess as the fruit fly (*drosophila*) of AI—that is, a simple system that allows the initial exploration of complex phenomena (Chase and 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-featuring simulation.

Although there has not been a lasting formal definition of AI,[5] chess computers do exemplify two of its core aspects. First, AI has generally been associated with tasks of a certain complexity, such as translation, hiring, and medical diagnosis (Agrawal, Gans, and Goldfarb, 2018; Brynjolfsson et al., 2019; Raj and Seamans, 2019). 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 abilities of even the most intelligent humans. Second, AI has been associated with solving complex tasks by emulating human thought (Simon, 1996). A human

[5] Definitions 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.

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 reflect that selective evaluation process based on codified heuristics (Simon and Schaeffer, 1992). Modern AI relies primarily on machine learning techniques where the AI follows rules derived from existing data. Despite these differences in the AI's backend, its defining purpose—mimicking the intelligence or behavioral pattern of humans—has remained the same.

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 chess games at tournaments. ChessBase includes meta-information on more than seven million games—such as each game's 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. Our player sample is defined by tournament activity and country of residence. All players in our sample have played tournament games.[6] We restrict our sample to players residing in either Western Europe or the (former) Soviet Union.[7] 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.[8]

---

[6] Tournaments usually require a minimum Elo rating to qualify; the lower bound in our sample is about 1,800.

[7] We exclude non-European countries of the Western world (Australia, Canada, the United States, etc.) to reduce the computational complexity of our data processing. Furthermore, tournament encounters with players from the (former) Soviet Union are less common for non-European players.

[8] The overwhelming majority of players in our data reside in Western Europe (see Appendix Figure A-1). 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.

We also create a dataset on chess computers. We first assemble an exhaustive list of over 500 commercial models with information on release date, strength, price, and manufacturer.[9] 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 (former) Soviet Union.

4.1.1 Research Design

We use the staggered diffusion of chess computers as exogenous variation in *access to chess computers*. Chess computers became available in Western Europe in 1977,[10] but not in the (former) Soviet Union until 1989.[11] This provides two discrete events that we exploit in our

---

[9] Our primary source for this is http://www.schach-computer.info/wiki, a community-based wiki. We cross-check and supplement this source with information on missing models and details from various online and offline sources.

[10] In 1977, the first commercial chess computer, *Chess Challenger,* went on sale in the West and quickly diffused among players. Several manufacturers then introduced stronger and cheaper models over the following years (see Appendix Figure A-2). Computer chess products became available in two forms: as dedicated chess machines (a microprocessor with a tailored software chess program) and as software running on standard PCs. Software did not play a significant role until the emergence of the Pentium processor in 1993. See Appendix Figure A-3 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 Appendix Figure A-4). In the same year, computer chess companies first topped 100 million USD in sales.

[11] Chess 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 and 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 and 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

analysis, allowing 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 (former) Soviet Union had access. This pattern should imply that any estimated differential effect of chess computers on player performance 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 and become more competitive in tournaments, we measure their progress in performance. We construct three distinct measures of a player's performance: (1) *Elo rating*—our primary measure—is used to quantify the relative skill of players in many zero-sum games, such as chess. A player's Elo rating increases (decreases) if the player wins (loses), with the increase (decrease) dependent on the difference in skill level.[12] Elo rating is well established in the chess community and is directly provided by our secondary data. It may, however, be subject to biases.[13]

We therefore use additional performance measures: (2) the *average centipawn loss*. The centipawn loss is a common measure to capture the quality of a single chess move. It is determined by comparing a player's advantage from playing a chosen move with the player's advantage from playing the best move available. To use this measure we need to assess all moves in our sample (21 million) and identify the optimal move in each case. To this end, we use Stockfish, a high-end

---

founded in 1989 and the leading Soviet chess magazine established a regular chess computer column in 1990 (Donskoy and Schaeffer, 1988).

[12] A novice player has an Elo rating of about 800, while that of the world chess champion is close to 2900.

[13] The Elo rating is comparative and therefore depends on the pool of players. In our empirical analysis, we compare the performance of Western and (former) Soviet players with each other. Differences in pool composition (for example, 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 Appendix Figure A-6) 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.

open-source chess engine.[14] 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.[15] The *average centipawn loss* is the player-year-specific average difference in advantage from all played moves.[16] As we illustrate in Appendix Figure A-5, 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.

Finally, as our third measure of performance we observe the likelihood of a player to win a game (*Game won*). For this measure, we focus exclusively on the outcome of direct confrontations between Western and (former) Soviet players who face off at tournaments despite political boundaries.[17] Taken together, we provide two performance measures at player-year level and one performance measure at game-level.

Independent variables. Our key independent variable, *Chess computer access*, captures the possibility of accessing chess computers. This variable is time-variant and indicates whether chess computers were available in the focal player's region. Western players have had access to chess computers since 1977, but Soviet players only since 1989. *Chess computer access* hence is a binary variable that equals 0 starting in 1970, equals 1 for Western players starting in 1977, and equals 1 for Soviet players starting in 1989. Note that our treatment variable captures whether the

[14] Stockfish has been consistently ranked first or second in chess-engine rating lists and is widely seen as the strongest conventional chess engine in the world. For more information, see https://stockfishchess.org/. Stockfish searches 20 moves deep and evaluates 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.

[15] We detail the calculation of average centipawn loss in Technical Appendix C.

[16] The centipawn is a chess-internal unit of measure to quantify advantage. A centipawn equals one-hundredth of a pawn (the least-valued figure in chess). Centipawns play no formal role in the game but prove helpful in evaluating positions and comparing possible moves.

[17] We thank the reviewer team for suggesting this operationalization.

focal player resided in a region with access to chess computers, not whether he or she actually used one. [18] 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.[19] 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:

$$y_{it} = \alpha_i + \beta_1 \, Chess\ computer\ access_{it} + age_{it} + \delta_t, \qquad (1)$$

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. Finally, because treatment might be correlated with the player's experience and seniority, we include age dummies to capture nonlinear life-cycle effects. Identification relies on the comparability of players from Western Europe and the (former) 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-featuring simulations in the form of chess computers helped chess players learn a form of strategic interaction. Table 1 shows the results concerning 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 annual Elo rating is

[18] We assume the same treatment intensity for all players in a region. Due to data limitations, we cannot factor in actual chess computer adoption and use at the player level.

[19] We run linear regression models with high-dimensional fixed effects, using the *reghdfe* Stata package as introduced by Correia (2015).

0.011 (Table 1, Column 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 eleven Elo points in the player's annual Elo rating. This increase is economically significant, making a treated player with an initial Elo rating of 2,000 about three percent more likely to beat an untreated player of previously equal rank under tournament conditions.

[INSERT TABLE 1 ABOUT HERE]

We test whether this effect is robust to our two alternative measures of player performance.[20] 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, Column 2). However, this effect is imprecisely estimated with a p-value below conventional thresholds of statistical significance. We also find consistent evidence for our third dependent variable: an increased likelihood that chess players with access to chess computers will beat those without it (see Table 1, Column 5).

We further run an event-study regression to alleviate concerns that chess players from Western Europe and the (former) Soviet Union are not comparable in skill development due to socio-economic differences between the regions. To this end, we regress player performance on a

[20] We 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 Appendix Table B-1, Columns 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 2,000 or higher (see Appendix Table B-1, Columns 5 and 6). Third, we restrict the sample to shorter time windows around the two discrete events in 1977 and 1989 (see Appendix Table B-2). All results correspond qualitatively and statistically with our main results.

full set of interaction terms between the binary indicator for Western Europe and the year fixed effects as independent variables.

[INSERT FIGURE 2 ABOUT HERE]

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 statistically significant 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).[21]

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 regional instead of 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" than "white".[22] Consequently, players training with chess computers typically played "white." In tournaments, however, a player is as likely to

[21] Moreover, we find that the magnitude of the point estimates increases during the 1980s. This effect could be driven by the ongoing diffusion of chess computers within Western Europe, but it may also be that the treatment effect itself becomes stronger due to the evolution of chess computers. That is, chess computers became a treatment for a larger part of the population along the skill distribution and this part of the population also benefited more.

[22] Chess 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.

play "black" as "white." If the suggested improvement is indeed due to training with chess computers, one would expect that players 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 and statistically more significant for "white" games" than for "black" games. Whereas the statistical test of the difference between these two coefficients is not significant by conventional thresholds (with a p-value between 0.1 and 0.2), we deem the difference nonetheless noticeable.[23] This non-uniform effect rules out many non-technological explanations for the improved performance of players in regions with access to chess computers.

Taken together, our results illustrate that AI-featuring simulations in the form of chess computers help people learn how to engage in strategic interactions. These quantitative results are buttressed by personal accounts of the disadvantage of lacking chess computers. Vladimir Kramnik, world chess champion 2006–2007, states:

> Today's player has access to [...] strong sparring partners whenever you want to play (and they never get tired); encyclopedic opening books [...]; comprehensive games collections [...]; and perfection in some endgames [...]. As a young chess player growing up in Russia and honing my skills, I had access to none of these tools. [...] I'm not complaining, just describing the not-so-distant past (Mueller and Schaeffer, 2018: 5).

In line with this, Knemeyer and Follet (2018) illustrate the virtue of chess computers for training:

> [They] can evaluate your play, calculate optimal moves, and annotate your game for you. The software can not only point out every mistake, it can also tell you the severity of that mistake and show you what the best move should have been. [...] [The chess computer] is like a really good coach, telling you, "You should have done this. You should have done that." [...] And that can be very useful for improvement, just studying and just seeing that (Knemeyer and Follet, 2018).

[23] This is underlined by the results of an additional test, where we regress chess computer access (as 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 and significant negative coefficient

4.1.5 Boundary Condition: AI Strength

We argue that decision-makers only benefit from training with AI if the AI can emulate opponents the trainee encounters in the field. This implies that chess players benefit from training with chess computers only if the chess computer can emulate opponents that are (at least) on par with the potential opponents at tournaments. To examine this, we leverage longitudinal variation in 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 played on grandmaster level and above. If only players whose skills are inferior to the chess computers at the time benefit from training with them, we expect to find heterogeneity in the treatment effect depending on whether or not the chess computers available at a given time matched a player's skill level.

[INSERT FIGURE 3 ABOUT HERE]

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* represents the Elo rating of the strongest commercially available chess computer in a given year, and *Chess computer superior* indicates whether the Elo rating of the strongest commercially available chess computer in a given year is higher than the focal player's Elo rating in that year. The intuition behind *Chess computer strength* is that trainees benefit the more intelligent per se their artificial training partner is. The intuition behind *Chess computer superior* is that only superior chess computers are helpful as they provide responses that the players cannot anticipate.

The full model is:

$$y_{it} = \alpha_i + \beta_1 \ Chess\ computer\ access_{it} + \beta_2 \ Chess\ computer\ access\ x\ Strength_{it}$$

$$+\beta_3\ Chess\ computer\ superior_{it}\ +\beta_4\ Chess\ computer\ access\ x\ Superior_{it}$$

$$+\beta_5\ Chess\ computer\ access\ x\ Superior\ x\ Strength_{it}\ +age_{it}\ +\ \delta_t, \quad (2)$$

where the dependent variable y is again a player's chess performance and the independent variables are *Chess computer access*, *Chess computer strength*, and *Chess computer superior* and their full set of interactions.[24]

We show that the magnitude of the treatment effect on Elo rating 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 and significant, suggesting that chess computer access had a stronger effect on the performance of players for 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 population) or quite 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. Columns 3 and 4 show the robustness of these interaction effects for our other performance measure at year level—average centipawn loss.

[INSERT TABLE 2 ABOUT HERE]

These results suggest that a player benefits from access to chess computers predominantly if the latter provides a response that he or she could not anticipate. In other words, the worse the player, the more a chess computer can help.[25] Note that this heterogeneity also helps us strengthen

---

[24] We 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 (former) 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.

[25] To 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 A-8). If we overstate the year-specific strength by 250 (500) Elo points, the coefficients of the independent variable Chess computer access x Chess computer superior x Chess computer strength

the identification of the main effect of chess computers, as we examine within-population variation.[26]

In summary, this finding of a heterogeneous treatment effect by chess computer strength is (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: the positive effect of AI on learning is confined to instances where the AI can provide sufficiently intelligent responses.

4.1.6 Strategic Implication: Catch-up and Convergence

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

We test 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 considerably *smaller* for Western players than for Soviet players (Figure 6). This suggests that chess computers foster skill convergence and intensify the competition.

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 mostly remain significant. These findings underscore the skill-level–dependent heterogeneity in treatment, as we would expect if only superior chess computers provide effective training opportunities.

[26] This 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 Appendix Table B-3).

4.2 Analysis - Part 2: Mechanism

4.2.1 Research Design

We argue 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 performance to depend on the degree to which a given player has access to human training partners (e.g., through friendly games in chess clubs). We would expect players constrained in their training opportunities to benefit more from artificial training partners. We therefore examine whether availability of other training opportunities moderates the effect of chess computer access on performance.

4.2.2 Variables and Econometric Model

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

$$y_{it} = \alpha_i + \beta_1\, Chess\ computer\ access_{it} + \beta_2\, Chess\ computer\ access\ x\ Chess\ event\ density_{it} + \beta_3\, Chess\ event\ density_{it} + age_{it} + \delta_t. \quad (4)$$

4.2.3 Results

We do indeed find that players with more training opportunities benefit less from access to chess computers (see Table 3). The interaction term in Column 1 (*Chess computer access x Chess event density*) suggests that a one-standard-deviation–higher density of events in the player's home country lowers the baseline treatment effect by about 25 percent. In Columns 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 x*

Chess computer superior x 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 valid substitute for human training partners as they provide them with training opportunities.

[INSERT TABLE 3 ABOUT HERE]

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

4.3: Analysis - Part 3: Imperfect Substitution

Given the evidence that artificial training partners are a substitute for human training partners in learning strategic interaction, we examine whether they constitute a *perfect* substitute for human training partners. We therefore examine 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: the consistency of performance. While human training partners frequently conduct idiosyncratic mistakes—so-called blunders—due to lack of concentration, limited memory or varying mood, the AI does not. Decision-makers who train with AI thus do *not* get exposed to the same amount of blunders. We would thus expect that players with access to AI are worse in 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 level of the game, where we choose as our performance measure the likelihood that a player wins in a direct encounter against a player from the other region (Table 4).[27] As a reference, we repeat the analysis from Section 4.2, where we find that access to chess computers has a positive effect on winning the game (Model 1). We further find that blunders made by the opponent, which we operationalize as moves that cause a centipawn loss in the 90th percentile, have a positive effect on winning the game (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 are less likely to benefit from these blunders compared to players without access.[28]

[INSERT TABLE 4 ABOUT HERE]

This finding suggests that training against highly rational artificial training partners does not prepare trainees equally well to exploit blunders conducted by boundedly rational human opponents. Given that an important part of strategic interactions is to exploit the opponents' blunders, AI remains to be limited in its ability to train decision-makers.

7 DISCUSSION

In our analysis of *whether* AI helps decision-makers learn, we find compelling evidence that chess players could improve their performance given access to chess computers. In particular, players with inferior skills and constrained training opportunities (thus, a competitive disadvantage) benefit particularly from and are able to catch up due to training with AI. Our analysis reveals the underlying mechanism—scalable artificial training partners resolve the

[27] The analysis on the game-level allows to establish a direct link between the occurrence of a blunder by the opponent and the outcome of the game. There is no reason to believe that the effect of blunders spills over to other games than the focal one.
[28] This result is robust to different operationalizations of blunders and to considering the number of blunders per game.

bottleneck of scarce human training partners. That said, artificial training partners remain imperfect substitutes for human training partners as they do not make idiosyncratic mistakes and thus keep decision-makers from learning how to exploit the blunders of human opponents.

7.1 Learning: AI as Substitute for Human Training Partners

Our research contributes to one of the key questions in the learning literature: *what underlies the heterogeneity in learning progress* (Darr, Argote, and Epple, 1995; Argote, 2012; KC, Staats, and Gino, 2013; Anand, Mullite, and Ren, 2016). Studying decision-makers' efforts to learn strategic interactions, we illustrate that the scarcity of qualified human training partners often constitutes a bottleneck.

Our study also illustrates how the bottleneck of scarce human training partners can be overcome thanks to artificial training partners, which provide responses comparable to those trainees receive from human training partners. Whether and where AI will provide meaningful substitutes to human training partners in other contexts than the one we study remains to be seen. However, if AI does, it will be highly consequential by "democratizing" responsive training of strategic interactions (cf. Bloom, 1984). As a matter of fact, the recent development of powerful chatbots (Mollick, 2022) and successfully negotiating AI (Kramar et al., 2022) suggests that AI can emulate training partners for a wide range of interactions. Finally, our study suggests that the learning outcomes of training with artificial instead of human training partners differ in potentially important dimensions. Specifically, the exposure to boundedly rational humans and their flawed decision-making remains important to learn how to identify and exploit idiosyncratic mistakes by the competitor. Where AI cannot realistically mimic the bounded rationality of humans, humans do not become superfluous in training strategic interactions. The testbed provided by AI-featuring simulations to train people may thus ill-prepare them on some dimensions.

7.2 AI in Management and Strategy: Training People

A key question for researchers and practitioners is how artificial intelligence (AI) will affect management and strategy (Raj and Seamans, 2019; Choudhury, Starr, and Agarwal, 2020; Iansiti and Lakhani, 2020; Tong, Jia, Luo, and Fang, 2021; Allen and Choudhury, 2022). Research has suggested that AI displaces or complements humans in the fulfillment of various tasks (Brynjolfsson and Mitchell, 2017; Acemoglu and Restrepo, 2018; Tschang and Almirall, 2021). We show that AI can also help decision-makers learn—a so-far neglected implication. The insight that technologies can help decision-makers learn is not new (Argote, 2012; Edery and Mollick, 2008; Iyengar, Sweeney, and Montealegre, 2015). What is new about learning with the help of AI is that it allows the simulation of human interactions in which decision-makers receive intelligent responses to their actions.

We illustrate that AI can substitute for human training partners, thereby resolving the trade-off between availability and responsiveness that defines the landscape of conventional training modes. AI-featuring simulation do so by offering the experiential learning that is comparable to training with (responsive) human training partners, and by being scalable and thus highly available (as other technology-based training modes). Demarcating AI-featuring simulation from other training methods also illustrates why it can be beneficial in a world where other training modes already exist.

We also illustrate that the relevance of AI for training purposes depends on the state of evolution of the AI and the skill-level of the trainee. Our results from chess computers, a field in which AI arrived early, suggest that weaker actors are the ones who may benefit first from AI in training.[29] However, the evolution of chess computers also showcases the tremendous progress of AI to the point where it eventually gains the potential to train even the very best.

We further contribute to the recently emerging stream of empirical research on AI by providing a large-scale field level historical dataset. Thanks to a common measure of performance in chess, we can also make progress in AI more tangible by comparing it to the performance of humans from the amateur to the world champion level. Research on technology and strategy has illustrated the value of documenting how a technology's performance evolves (Christensen, 1992; Adner and Kapoor, 2016), but such documentation has been missing for AI.

Our research also informs how AI-featuring computer simulations may be designed to train people even more effectively. Oddly enough, AI may be a more effective training tool if it showed from time to time human-like flaws, such as blunders. It would keep the trainee on constant alert to look for such mistakes and learn how to exploit them.

7.3 Competitive Strategy: Origin and dynamics of competitive advantage

Our analysis informs research on the origin and dynamics of competitive advantage. We illustrate that access to human training partners can constitute a source of competitive advantage for decision-makers. Prior research that examines the role of people the focal actor has access to, has typically emphasized their role as collaborateurs, trading partners, or sources of information (Burt, 2005; Uzzi and Spiro, 2005); our emphasis on training partners points to a so far neglected function. Our findings also showcase how AI can undermine the competitive advantage coming

---

[29] The question of who benefits from AI-featuring simulations may also have a bearing on their diffusion. Intuitively, one may expect that AI, long associated with innovation leaders, would be adopted first by people with greater skills and resources. Our analysis suggests the opposite: AI may in fact first be adopted by those with lower skills—as they can benefit from it even in its infant stage—and by those with fewer resources—as those with greater resources can afford the human training partners the AI is supposed to substitute. By implication, one may argue that grammar checkers in word processing software are welcomed first and foremost by foreign speakers who cannot afford copy editors.

from access to many human training partners: AI-featuring simulation allowing anyone to train interactively, eroding such a positional advantage.

Second, we complement prior research that has examined population-level distribution in performance (Lenox, Rockart, and Lewin, 2006, 2010; Porter and Siggelkow, 2008; Kapoor and Agarwal, 2017). We found systematic variation indicating that *dis*advantaged actors, specifically players with inferior skills and with constrained training opportunities, benefit most from chess computers. Given that disadvantaged actors benefit earlier and more, AI-based training can lead to a more equal skill distribution—it levels the playing field. Research has documented enormous skill heterogeneity within industries (Lenox, Rockart, and Lewin, 2006, 2010). AI-featuring simulations can help address that because the disadvantaged actors in the left tail of the distribution are the most likely to benefit.

Our findings also inform research on competitive moves. Prior research has pointed to the importance of moves in strategic interactions (Ferrier, Smith, and Grimm, 1999; Smith, Ferrier, and Ndofor, 2005). We illustrate that AI-featuring simulations can play a key role in learning such competitive moves. Finally, we contribute to research on strategic interaction: we illustrate that decision-makers' ability to exploit blunders of their boundedly rational opponents constitute an important skill—a skill that is better acquired when training with boundedly humans than with consistently high-performing machines.

A key feature of our study is also that succeeding in strategic interaction implies to identify and successfully idiosyncratic mistakes by the competitor. While we suggest that AI-featuring simulations constitute a substitute to human training partners in the sense that they also provide intelligent responses, they are not perfectly human-like as the AI-featuring simulation is consistent in its performance and rarely conducts such blunders. Ultimately, our study shows that learning outcomes can differ between training with AI and training with human interaction partners, with relevance for their performance in competition with other humans.

7.4 Boundary Conditions

We illustrate that AI-featuring simulation plays a major role in the training of people in chess. While we do not claim this to be necessarily representative, we do believe that it is plausible that AI-featuring simulation will soon play a major role in training people in strategic interactions and - perhaps - other kinds of social interaction. In particular, we expect that the use of AI-featuring simulations becomes more likely where AI has a beneficial effect. This applies most likely to contexts where the following boundary conditions are fulfilled: (a) training creates a high return, (b) experiential learning through actual interactions is costly, and (c) human training partners are scarce. For example, in the military, where strategic interactions are key to success, and where learning through actual interactions can cost lives, AI-featuring simulations are already employed.[30] We identified a further boundary condition that needs to be fulfilled: (d) AI-featuring simulations allow for a realistic emulation of human training partners. , as we show empirically, the AI must be good enough to emulate a human training partner the focal trainee is likely to encounter and whose responses cannot be easily anticipated. The realistic emulation of human training partners is particularly challenging in settings that are not as clearly defined and closed as the setting of chess. It remains open if (or at least when) AI will be able to emulate human behavior in a complex real life setting.

The requirement that the AI needs to be sufficiently strong to emulate human training partners includes another implication: the AI will in theory be able to replace the focal human.

[30] For 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 vs. 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.

Based on this argument, we expect that training with AI will be most relevant in occupations and activities in which human involvement remains socially desired (Glikson and Woolley, 2020), legally required (Tambe, Cappelli, and Yakubovich, 2019), or simply unavoidable (Brynjolfsson and McAfee, 2012). For example, even if AI could provide better diagnosis than a doctor can, people might still prefer to be diagnosed by a doctor.

Deploying AI-featuring simulations to train decision-makers may also be bounded if the AI has algorithmic biases or complicates causal inference. Current AI algorithms are not hard-coded; they emerge from patterns in data. However, AI algorithms based on selective or biased data may aggravate human biases (Lambrecht and Tucker, 2018; Cowgill, 2019; Choudhury, Starr, and Agarwal, 2020). AI-featuring simulation may then pass on such biases to trainees.[31] It is worthwhile to note that these biases inherited from humans remain systematic and do not create idiosyncratic mistakes that turn out to be crucial to approximate the learning outcome of training with humans. Beyond the outlined biases, causal ambiguity may bound the learning effect of AI-featuring simulations. That is, if trainees cannot figure out why the AI responds as it does (cf. Agrawal, Gans, and Goldfarb, 2018; Puranam, Shrestha, He, and von Krogh, 2020), its role in training them and its ability to replace human training partners and coaches would be bounded.

7.5 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 due to 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

[31] Note that the AI in chess computers is primarily based on predefined algorithms. Chess computers are therefore not prone to behavioral biases, but neither can they make use of them as human players sometimes can. An action may be best conditional on the other player responding as predicted (by the computer), but may then be dominated by an unexpected move.

abundant anecdotal evidence of chess computer usage, we do not know which players used what kinds of chess computers to what extent.[32] This limitation most likely results in an *under*estimation 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 semi-professional tournament players, not of amateurs. We would, however, expect—in line with our results—that amateurs benefited even more and earlier from chess computers, as the computers exceeded those players' skill level early on.

7.6 Managerial Implications and Use Cases

We propose that our examination of the effect of AI in the realm of chess provides a glimpse into the future and that training with AI-featuring simulations will play an increasing role in training decision-makers in strategic interactions—or, more broadly, in training people in all kinds of social interactions. Our general argument has been that AI-featuring simulations provide training similar to training with human training partners or coaches. We thus see potential for AI-featuring simulations in virtually all domains in which people benefit—or at least could benefit—from human training partners. Specifically, we expect use cases of AI-featuring simulations to emerge where a high demand for human training (i.e., where training with intelligent feedback is beneficial) cannot be met because such training opportunities are scarce or too costly.

Training with AI-featuring simulations may help managers to improve their ability to engage in *strategic interactions with other firms*. Imagine, for example, that a firm participates in a multi-stage auction (e.g., for mobile bandwidth licenses). In such a (closed) setting, the firm may

[32] Appendix Figure A-7 lists the main findings of a chess computer market survey of West German chess 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.

rely on AI-featuring simulation to prepare for participation and to test-run bidding strategies. The same may be true for companies involved in complex strategic interactions. For example, airlines compete with one another by taking various actions with respect to plane capacity, route planning, prices, frequent flier programs, and so on. Such strategic interactions are difficult to manage and airlines often find themselves in price wars that they have triggered accidentally (Miller and Chen, 1994; Gimeno, 1999). Airline executives who can train with AI-featuring simulation may become more savvy navigators of such interactions. In fact, firms already hold "war games" to test how certain strategies may play out (Courtney, Koller, and Lovallo, 2020), but such training—typically involving coaches or conventional simulations—is either far more expensive or less realistic than AI-featuring simulations would be.

Training with AI-featuring simulations may also help managers to improve their ability to engage social interactions with people. The presence of social skills has been linked to employee and firm performance (Dimitriadis and Koning, 2021), but training in social interactions is expensive due to the need for human training partners or coaches. We suggest that AI-featuring simulations could provide, for instance, conversation training. In an informal interview, the head of human resources of a service company with 700,000 employees pointed out to us that they were considering using AI to provide more cost-efficient customer service training for employees. A prime use case he saw was that employees in call centers, rather than accumulating experience with real customers, could first try out conversational strategies with AI-featuring chatbots. Generally, organizations can use AI-featuring simulations to provide interactive personalized training to many more employees than would have been possible with expensive coaches.

8 CONCLUSION

Using the case study of chess computers, we examine the linkage between AI-featuring simulations and learning strategic interactions—whether people benefit, who benefits, why people benefit, and what they learn. It is plausible that AI will democratize skill development with respect to learning social interactions. While accumulating experience in most interactive tasks still requires human training partners, AI-featuring simulations could soon provide worthwhile training opportunities at considerably lower cost. 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.

REFERENCES

TBA

Figure 1: Schematic overview of training methods for strategic interactions

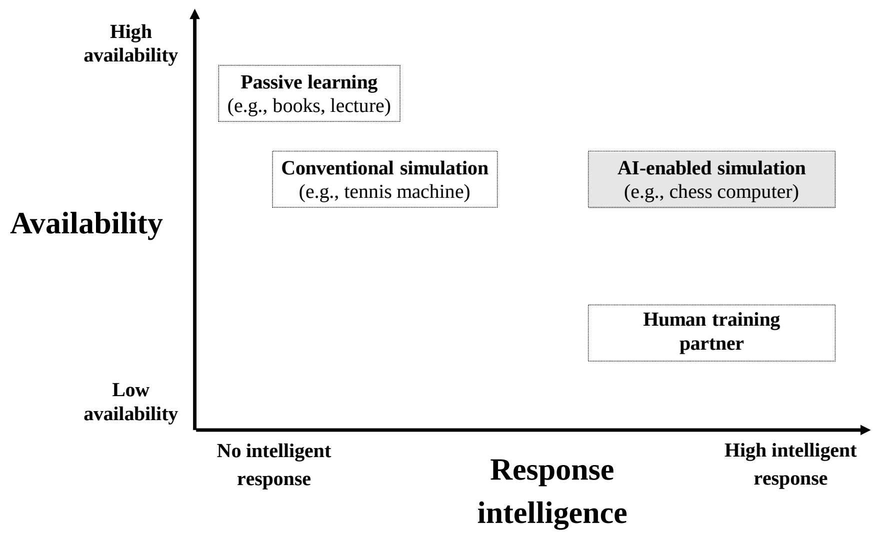

Figure 2: Event study: the effect of region (Western Europe) on chess performance (Elo rating)

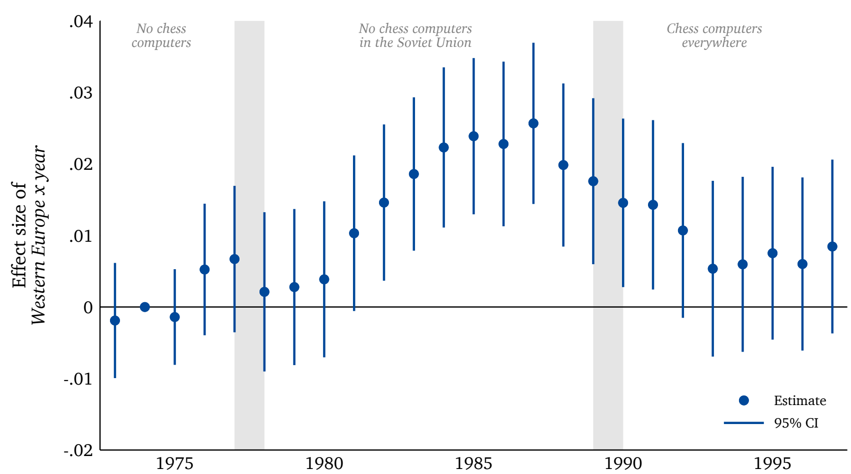

Notes: This figure plots the point estimates and 95% confidence intervals of the effect of *Western Europe x Year* on Elo rating. The regression specification is as follows: $y_{it} = \alpha_i + Western\ Europe_j \times \sum_{k=1970;k\neq -1973}^{2000} \beta_k + age_{it} + \delta_t$.

Figure 3: Chess computer strength over time

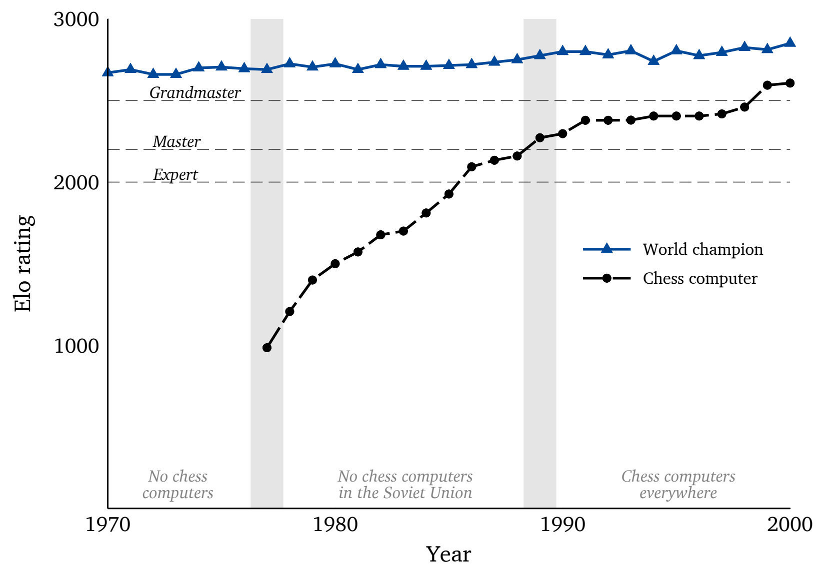

Notes: This figure illustrates the progress in chess computer strength over time in Elo scores. For each year, the upper bound strength of all available commercial chess computers is considered. Non-commercial chess computers (e.g., Deep Blue) are excluded.

Figure 4: Skill distribution over time and region

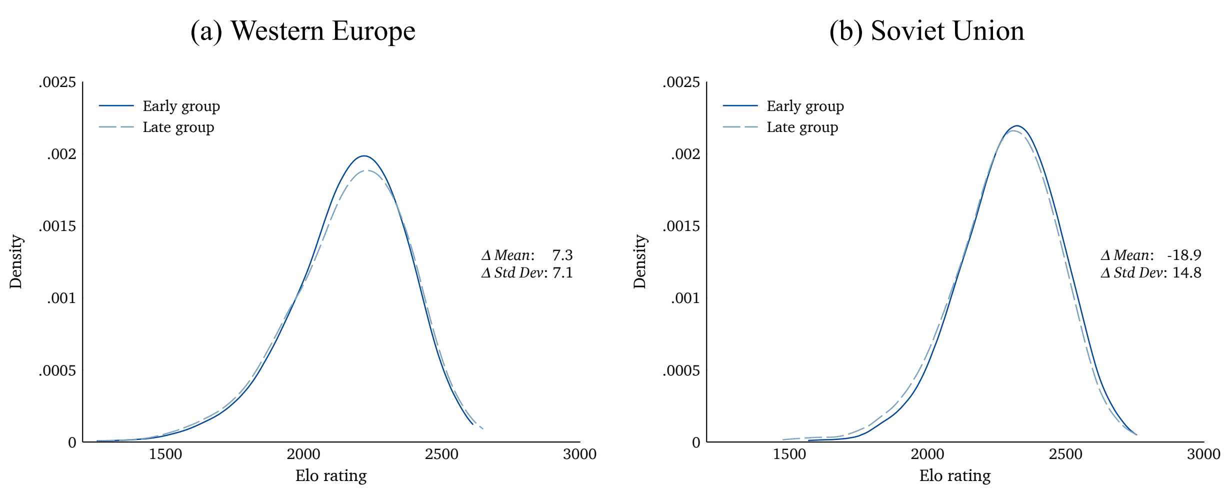

Notes: These two figures plot the skill distribution as kernel densities for Western and Soviet chess players active in 1980 (early group) and active in 1988 (late group). Figure 4a depicts the skill distribution of active players from Western Europe. Figure 4b depicts the skill distribution of active players from the Soviet Union. Δ*Mean* represents the difference in the mean Elo rating between the early group and the late group. Δ*Std Dev* represents the difference in the standard deviation of Elo rating between the early group and the late group.

Table 1: Chess performance over time

Sample: All players	(1)	(2)	(3)	(4)	(5)
DV:	Elo rating	Average centipawn loss			Game won
		All	White	Black	
Chess computer access	0.011***	−0.362	−0.706**	−0.381	0.060***
	(0.003)	(0.302)	(0.342)	(0.440)	(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

Notes: 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 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 player level in parentheses.

Table 2: Chess performance over time – heterogeneity by chess computer strength

Sample: All players	(1)	(2)	(3)	(4)
DV:	Elo rating		Average centipawn loss	
Chess computer (CC) access	0.015***		−0.050	
	(0.003)		(0.368)	
× CC strength	0.016***		1.084*	
	(0.005)		(0.601)	
× CC superior	0.038***	0.040***	0.074	0.217
	(0.009)	(0.010)	(0.650)	(0.673)
× CC superior × CC strength	0.036***	0.037***	−1.896**	−2.169***
	(0.009)	(0.009)	(0.766)	(0.771)
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

Notes: 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. Coefficients for *Chess computer (CC) superior* omitted. Robust standard errors clustered at player level in parentheses.

Table 3: Chess performance over time – mechanism: training opportunities

	(1)	(2)	(3)	(4)
DV:	Elo rating	Elo rating		Tournament attendance
Sample:	All	Chess event density: Low	High	All
Chess computer (CC) access	0.012***			
	(0.003)			
× Event density	−0.003***			
	(0.001)			
× CC superior		0.040***	0.040***	−3.110***
		(0.010)	(0.010)	(0.519)
× CC superior × CC strength		0.041***	0.027**	−8.323***
		(0.012)	(0.011)	(0.866)
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

Notes: Columns (1) to (6) show the estimates of linear model regressions with high-dimensional fixed effects. Coefficient for *Chess event density* in Column (1) and coefficients for *Chess computer (CC) superior* in Column (2) to (6) omitted. 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 player level in parentheses.

Table 4: Chess performance in case of blunder

Sample: All players	(1)	(2)	(3)
DV:		Game won	
Chess computer access	0.060***	0.057***	0.077***
	(0.014)	(0.011)	(0.013)
Blunder by opponent		0.299***	0.324***
		(0.006)	(0.013)
Chess computer 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

Notes: 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.

ONLINE APPENDIX

A Appendix: Figures

Figure A-1: Human chess players over time and region

(a) Share of players by region

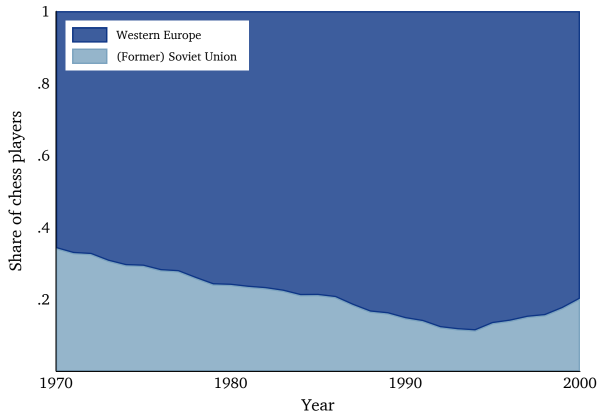

(b) Number of players by region

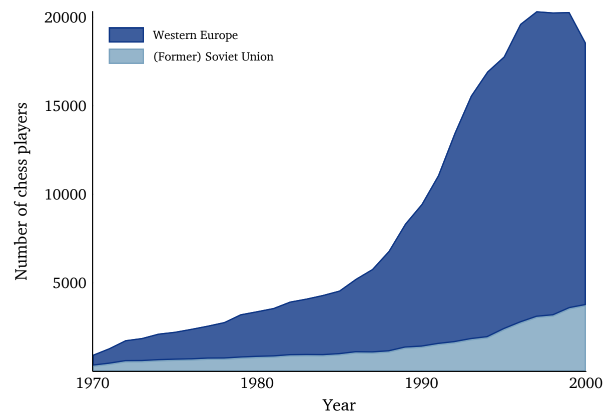

(c) Activity over time by player cohort

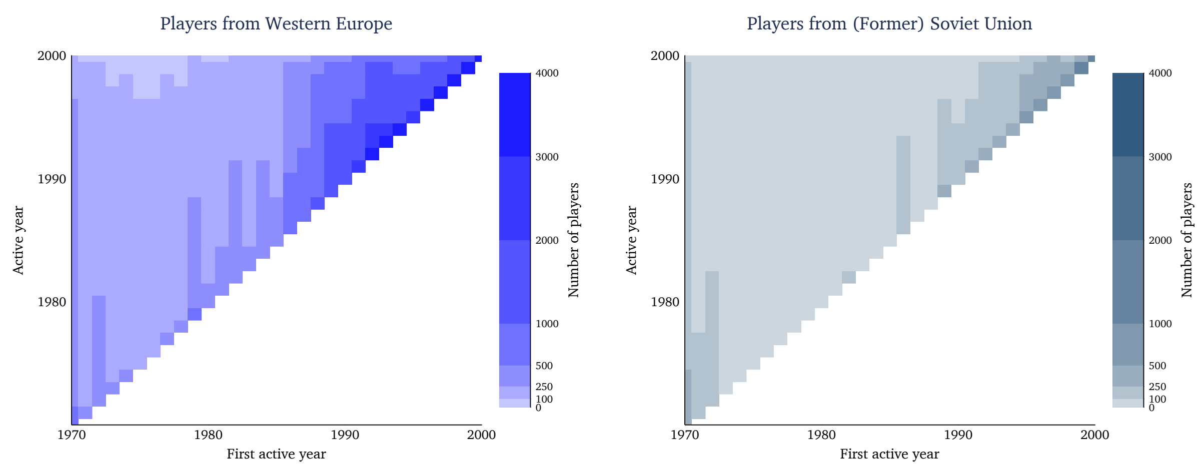

Notes: This set of figures illustrates the composition of our panel of human chess players over time and region. Figure A-1a depicts the share of active players from Western Europe and the (former) Soviet Union over time. Figure A-1b depicts the absolute count of active players from Western Europe and the (former) Soviet Union over time. Figure A-1c depicts for each region separately the number of active players over time by their first year of activity.

Figure A-2: Chess computer models and manufacturers over time and region

(a) New models over time

(b) Manufacturers over time

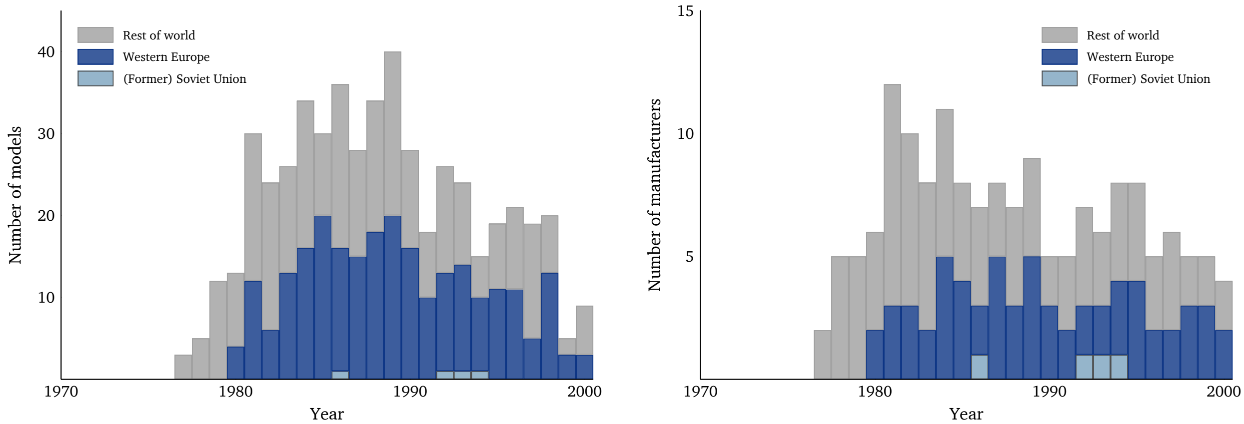

Notes: These two figures illustrate the number of chess computer models (Figure A-2a) and chess computer manufacturers (Figure A-2b) over time and region.

Figure A-3: The first personal (dedicated) chess computer *Challenger 1* (1977)

Notes: This figure shows a picture of the first personal chess computer: the *Challenger 1* by *Fidelity Electronics*.

Figure A-4: USCF promotion of a chess computer for members (1982)

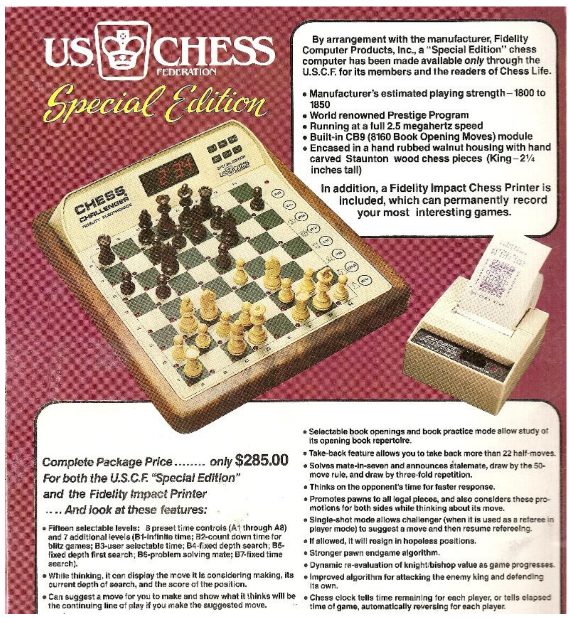

Notes: This figure shows a personal chess computer model advertised by the United States Chess Federation (USCF) to its members.

Figure A-5: Relationship between Elo rating and average centipawn loss

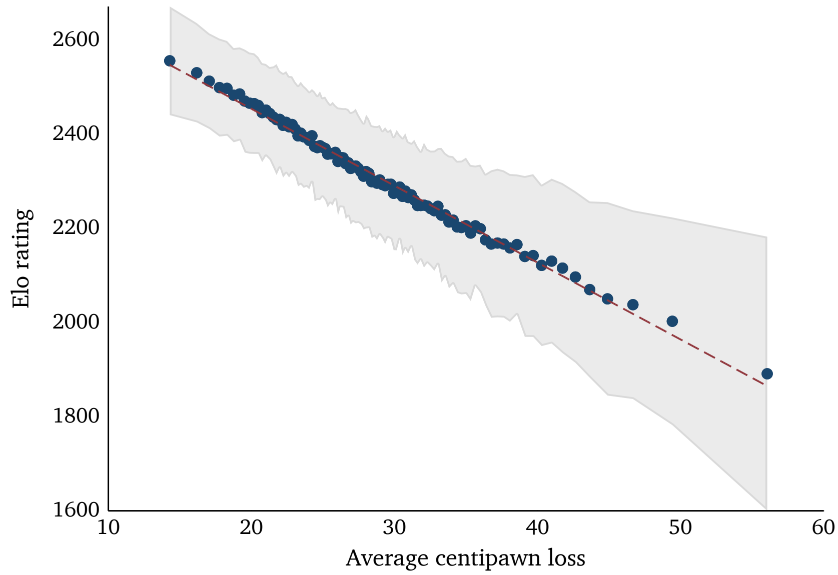

Notes: This figure plots the mean ELO rating and the mean average centipawn loss for 100 bins. The dashed line depicts the fit line. The grey area indicates the band of one standard deviation above and below the mean in each bin. The disaggregated level of observation is the human chess player by year. Average centipawn losses based on fewer than 250 moves per year excluded.

Figure A-6: Home and away games of human chess players by region

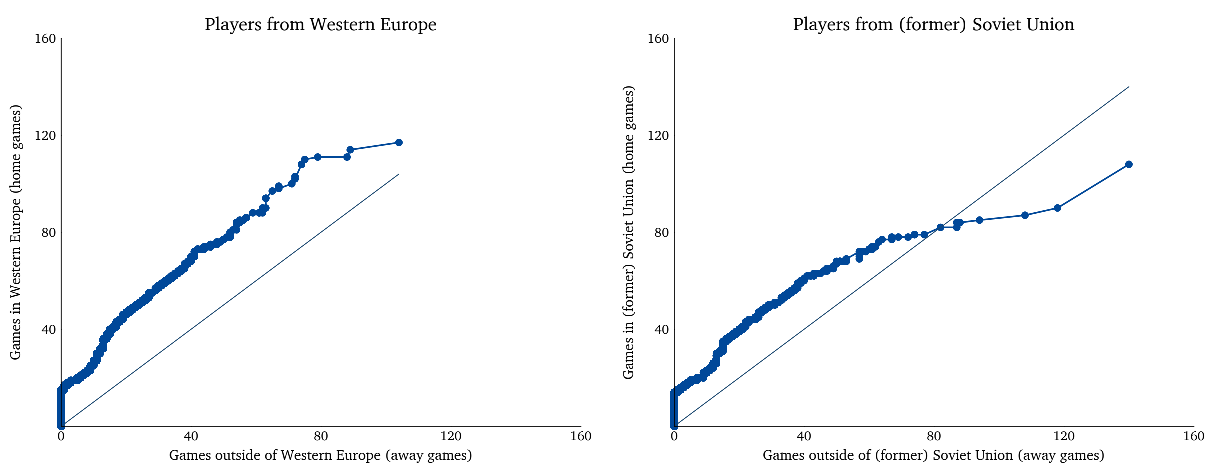

Notes: This figure illustrates the distribution of home and away tournament games of human chess players by each region separately. The figure is a quantile-quantile plot that compares the distributions of tournament games in the domestic region (Western Europe and the (former) Soviet Union, respectively) and games outside of it. Points above the 45 degrees line indicate a larger share of games played in the player's domestic region. Only tournament games before 1989 considered.

Figure A-7: Count of chess masters and grandmasters over time and region

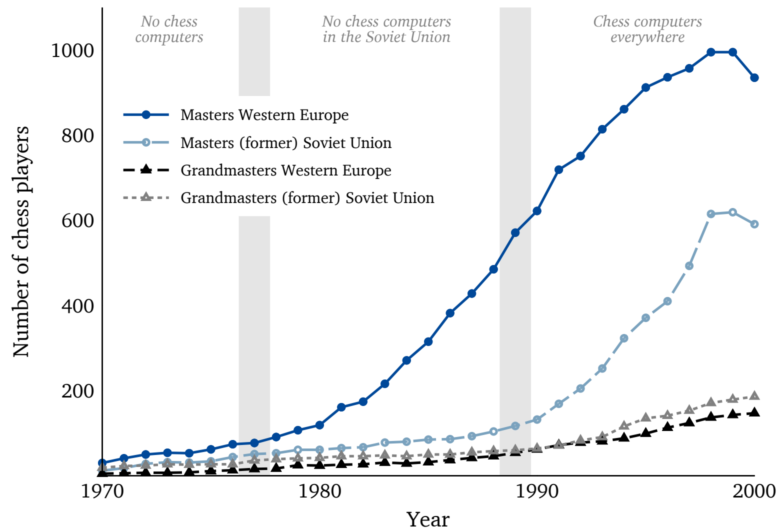

Notes: This figure depicts the annual counts of active human chess players with a FIDE master or grandmaster title by region. In general, a chess master title is achieved with an Elo rating of 2,200 and above. A grandmaster title is achieved with a rating of 2,500 and above. Once awarded, the player holds the title for life.

Figure A-8: Estimates by manipulated chess computer strength (placebo regressions)

(a) Elo rating

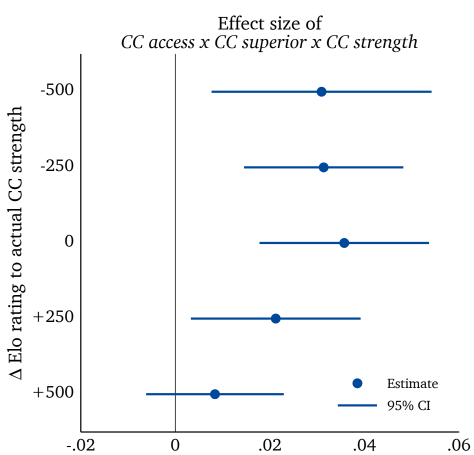

(b) Average centipawn loss

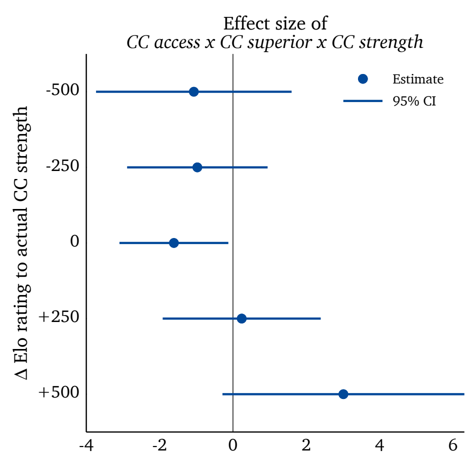

Notes: These two figures provide a comparison of estimates based on regressions with differently manipulated chess computer strength. Figure A-8a states the effect size of *Chess computer (CC) access x Chess computer (CC) superior x Chess computer (CC) strength* on player performance (Elo rating). Figure A-8b states the effect size of *Chess computer (CC) access x Chess computer (CC) superior x Chess computer (CC) strength* on player performance (average centipawn loss). The estimates are based on regressions as specified in equation (2). In each model, *Chess computer strength* is manipulated by adding (subtracting) 250 or 500 Elo points to (from) the chess computer's actual rating. Note that the variables *Chess computer access* × *Chess computer strength*, *Chess computer superior*, and *Chess computer access* × *Chess computer superior* are a function of the manipulated computer strength.

Figure A-9: Results of a 1987 market survey on West German chess players (in millions)

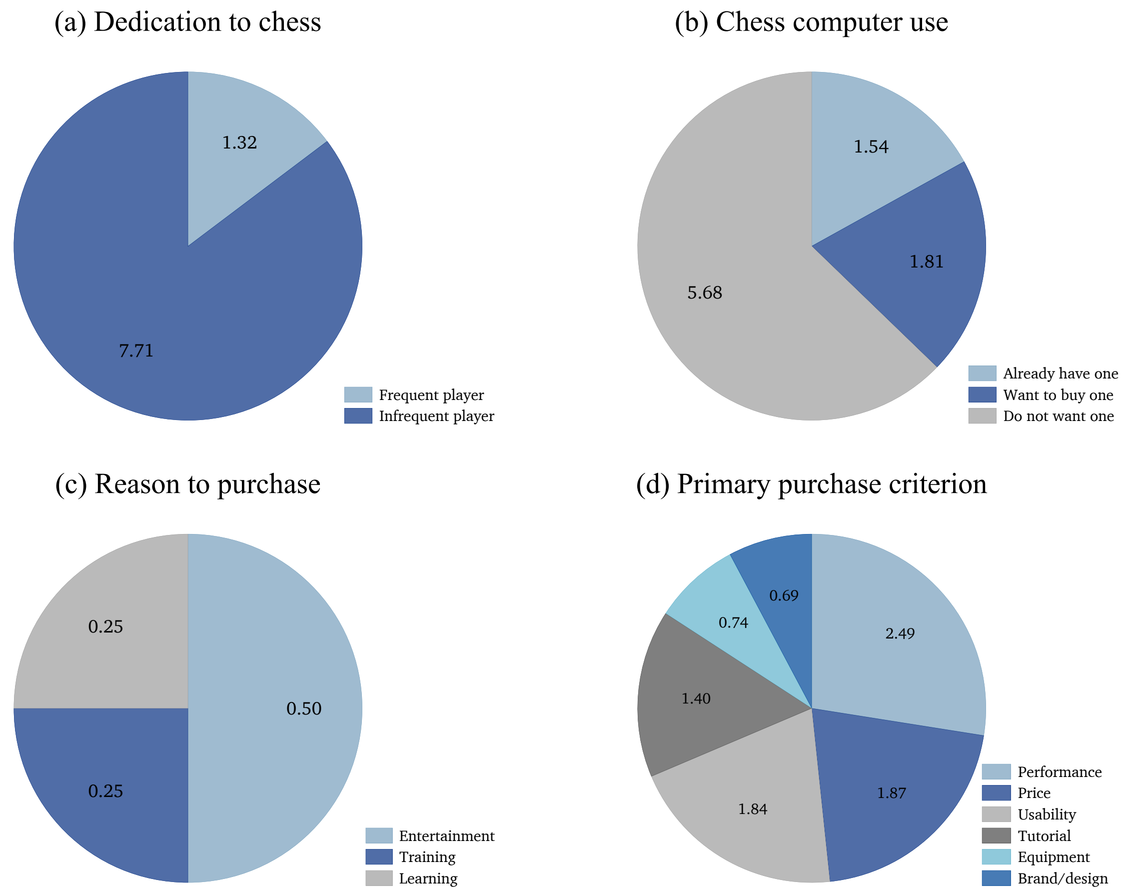

Notes: Market analysis on behalf of chess computer manufacturer *Hegener+Glaser Aktiengesellschaft*. Survey conducted between June and August 1987. Results based on 2,014 respondents. The survey estimates a population of about 9 million chess players in West Germany.

B Appendix: Tables

Table B-1: Chess performance over time for different player subsamples (Elo rating)

Sample:	(1)	(2)	(3)	(4)	(5)	(6)
All players with:	≥ 90 games		≥ 9 years active		≥ 2000 Elo	
	Elo rating	Elo rating	Elo rating	Elo rating	Elo rating	Elo rating
Chess Computer (CC) access	0.013	0.023	0.011	0.019	0.011	0.017
	(0.004)	(0.004)	(0.003)	(0.003)	(0.002)	(0.002)
× CC strength		0.023		0.020		0.015
		(0.007)		(0.004)		(0.003)
Player FE	Yes	Yes	Yes	Yes	Yes	Yes
Player age FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	52679	52679	108537	108537	128054	128054
Players	3355	3355	6483	6483	13609	13609
Log-likelihood	67130	67149	149666	149700	211545	211574

Notes: Columns (1) to (6) 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.

Table B-2: Chess performance over time for restricted observation windows (Elo rating)

Sample:	(1)	(2)	(3)	(4)	(5)	(6)
All players in	1970-1988		1977-2000		1985-1995	
	Elo rating	Elo rating	Elo rating	Elo rating	Elo rating	Elo rating
Chess Computer (CC) access	0.014	0.026	0.009	0.017	0.011	0.019
	(0.004)	(0.005)	(0.003)	(0.003)	(0.003)	(0.003)
× CC strength		0.022		0.020		0.020
		(0.004)		(0.004)		(0.004)
Player FE	Yes	Yes	Yes	Yes	Yes	Yes
Player age FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	43210	43210	164993	164993	172944	172944
Players	4619	4619	19343	19343	19463	19463
Log-likelihood	65217	65267	222503	222537	231729	231761

Notes: Columns (1) to (6) 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.

Table B-3: Chess performance over time by each region separately (Elo rating)

Sample:	(1) Soviet players	(2)	(3) Western players	(4)
	Elo rating	Average centipawn loss	Elo rating	Average centipawn loss
Chess computer (CC) access × CC superior	−0.026	−0.321	−0.014	−0.403
	(0.005)	(0.570)	(0.002)	(0.178)
Chess computer (CC) access × CC superior × CC strength	0.054	−2.286	0.042	−1.586
	(0.020)	(2.330)	(0.010)	(0.808)
Player FE	Yes	Yes	Yes	Yes
Player age FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	30854	33948	142090	151032
Players	3458	3914	16005	17628
Log-likelihood	41322	−117961	190766	−555417

Notes: 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.

Table B-4: Learning curve in chess over time (Δ Elo rating)

Sample: All players DV:	(1)	(2) Δ Elo rating	(3)
Chess computer (CC) access	0.004	0.008	
	(0.001)	(0.002)	
× CC strength		−0.003	
		(0.002)	
× CC superior		−0.005	−0.001
		(0.005)	(0.005)
× CC superior × CC strength		0.070	0.072
		(0.008)	(0.008)
Player FE	Yes	Yes	Yes
Player age FE	Yes	Yes	Yes
Year FE	Yes	Yes	No
Region-Year FE	No	No	Yes
Observations	151070	151070	151070
Players	17052	17052	17052
Log-likelihood	188707	189437	189490

Notes: Columns (1) to (3) 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.

C Technical Appendix: Evaluating Moves

We use the *python-chess* module to implement an evaluation algorithm and *Stockfish 10* as chess engine to evaluate performance of chess players at the move level. The resulting *average centipawn loss* measure follows the definition of Guid and Brakto (2011). We first give a motivation for why we use this measure and then detail its construction.

Background

The measure introduced by Guid and Brakto (2011) allows for the comparison of chess players even if they never directly played against each other. The lack of encounters may be due to local clustering, distance, or because their timespans of being an active player do not overlap. This property holds because the evaluation is independent of the overall quality of play of the environment in which the player to be evaluated interacts and is solely based on move-level performance. By contrast, it does not hold for the common Elo rating, a cumulative game-level based measure, for which the actual outcome of the game is compared to the predicted outcome, which in turn is determined by the players' difference in ratings. In other words, when the lower rated player wins, more points will be obtained from his opponent then when the higher rated player does. A player's change in Elo thus depends on his opponent's, as well as his own current point-level. Elo therefore is only a relative skill measure that ignores differences in the means of populations, ruling out comparisons of players across different environments. Since the data we use spans almost three decades and multiple countries on different continents, Elo is not suitable for our analysis.

The measure by Guid and Brakto (2011) allows for absolute comparison because, unlike Elo, it uses an external evaluator that is not part of the pool of players to be evaluated and that is stable across locations and times. In their "Computer analysis of World Chess champions", Guid and Brakto (2006) use chess-engine *Crafty* to quantify the advantage of each player over its opponent in the position preceding and the one following a move. After the difference in advantage that results from the move is calculated, this number will be compared to the evaluation of the hypothetical, best possible move as suggested by the engine. The resulting move-level skill-measure can be interpreted as the number of advantage-points a player misses to those of a chess-computer that was confronted with the exact same position.

One point of concern about this measure is that the engine could discriminate against certain player types, independent of the strength of the moves played. Guid and Brakto (2006) mention that positional players might on average face less complex positions than tactical players. The authors therefore developed a measure of position complexity. Their metric is built on the notion that, while evaluating the same position under different search depths, the engine's suggestion for the best possible move might change with an increase in depth. The more that happens, the more this position is seen as being more complex than another position in which the suggested best possible move is invariant to changes in depth. When using a low search depth already lets the engine detect the "true" optimal move, the position can be seen as being less complex.

For the analysis, we chose Stockfish 10, which is sometimes considered the strongest chess engine available (Wilkenfeld, 2019).The analysis excludes the opening game, which is based mostly on theory and still challenges current chess programs in delivering accurate evaluations (Guid and Brakto, 2006). Evaluation in each game therefore starts at the 12th move of the white player.

Method

Our initial dataset contains game-level observations about all the moves played. The algorithm will loop through all games $g \in \{1, 2, \ldots n\}$ and all halfmoves of game g that have indices $i \in \{1, 2, \ldots, m_g\}$. Each single move of a game is called a halfmove or ply, whereas a fullmove always contains two moves, one by the white and one by the black player. We only use the notion of a halfmove in the subsequent chapters.

As a basis, we define $position_{ig}$ as a board-position before halfmove number i of game g is played. $evaluation_{igp}$ represents the chess-engine's evaluation of $position_{ig}$ from perspective of player $p \in \{0, 1\}$, where:

$$evaluation_{igp} = f_p(position_{ig}, \theta).$$

The function f_p represents stockfish's internal evaluation algorithm that takes a board-position as an input and returns a score, representing the advantage from perspective of player p over its opponent. θ contains engine parameters such as the chosen search depth δ and contempt factor κ. The type of score the engine will return depends on whether there are end-nodes in the current search-tree. If there are, the return value will be a mate-score, representing the number of moves until the player is expected to mate his opponent in case of a positive value, or the number of moves until he is expected to be mated in case of negative values. The usual case is when there are no end-nodes present. The engine will then return a centipawn score that represents the advantage of a player measured as a hundredth of a pawn. Since we want to calculate differences in advantage, we need to use the same units and therefore use another function to convert mate- into centipawn-scores:

$$\begin{aligned} advantage_{igp} &= g(f_p(position_{ig}, \theta)) \\ &= g(evaluation_{igp}), \end{aligned}$$

where

$$g(evaluation_{igp}) = \begin{cases} 32768 - evaluation_{igp} * 1000, & \text{if} \quad evaluation_{igp} \in \mathbb{Z} \setminus \{0\} \times \{MATE\} \\ evaluation_{igp}, & \text{if} \quad evaluation_{igp} \in \mathbb{Z} \times \{CP\} \end{cases}$$

The constant in this linear function is the maximum value of a 16-bit integer measured in centipawns. We will come back to why the exact number is of less importance. For each mate-score in $evaluation_{igp}$, we substract 1000 centipawns from this number, representing the willingness to sacrifice the queen (which is worth 900 CP) in order to delay a checkmate by one move. The engine also returns a suggestion for the best possible move to play in the current board position.

In order to proceed, lets define m_{ig} as a move and part of the space of all possible moves

$m_{ig} \in \mathscr{M}_{ig}$ in $position_{ig}$ that maps a given board position $\overline{position}_{ig}$ into $position_{i+1g}$:

$$\text{move } m_{ig} = \psi^{-1}|_{m_{ig}}(\overline{position}_{ig}, position_{i+1gp}), \text{with}$$

$$position_{i+1g} = \psi(\overline{position}_{ig}, m_{ig})$$

The best move $\widetilde{m}_{ig}$ is the maximizer of the engine's evaluation of $position_{i+1,g}$ for all possible moves $m_{ig} \in \mathscr{M}_{ig}$ in $\overline{position}_{ig}$ given engine parameters θ:

$$\widetilde{m}_{ig} = \underset{m_{ig} \in \mathscr{M}_{ig}}{\arg\max} f_p(\psi(\overline{position}_{ig}, m_{ig}), \theta)$$

The advantage of player p after playing this hypothetical move $\widetilde{m}_{ig}$ in $position_{ig}$ is:

$$\widetilde{advantage}_{i+1gp} = g(f_p(\psi(position_{ig}, \widetilde{m}_{ig}), \theta))$$

We then define a $loss_{igp}$ as being the evaluation of move m_{ig} according to Γ_p. This function simply takes the difference in advantage of player p before and after playing that move, given the engine parameters θ:

$$\begin{aligned}
loss_{igp} &= \Gamma_p(m_{ig}, \theta) \\
\Gamma_p(m_{ig}, \theta) &= g(f_p(position_{ig}, \theta)) - g(f_p(\psi(position_{ig}, m_{ig}), \theta)) \\
&= advantage_{igp} - advantage_{i+1gp}
\end{aligned}$$

The definitions above now allow us to define the move-level skill measure. It can be described as the evaluation of the move played but corrected by the evaluation of the best possible move:

$$\begin{aligned}
moveCPL_{igp} &= \Gamma_p(m_{ig}, \theta) - \Gamma_p(\widetilde{m}_{ig}, \theta) \\
&= \{g(f_p(\overline{position}_{ig}, \theta)) - g(f_p(\psi(\overline{position}_{ig}, m_{ig}), \theta))\} \\
&\quad - \{g(f_p(\overline{position}_{ig}, \theta)) - g(f_p(\psi(\overline{position}_{ig}, \widetilde{m}_{ig}), \theta))\} \\
&= g(f_p(\psi(\overline{position}_{ig}, \widetilde{m}_{ig}), \theta)) - g(f_p(\psi(\overline{position}_{ig}, m_{ig}), \theta))
\end{aligned}$$

The final skill-measure of player p in game g is calculated by aggregating the move-level evaluations:

$$\begin{aligned} aCPL_{p,g} &= \sum_{i=0}^{m_g} \Phi(moveCPL_{igp}) * d_{i,g_p} \\ &= \sum_{i=0}^{m_g} \Phi(\Gamma_p(m_{ig},\theta) - \Gamma_p(\tilde{m}_{ig},\theta)) * d_{ig}, \end{aligned}$$

where

$$d_{igp} = \begin{cases} 1, & \text{if } \Gamma_p(m_{ig},\theta) \in [-200,200] \wedge \Gamma_p(\tilde{m}_{ig},\theta) \in [-200,200] \\ 0, & \text{otherwise} \end{cases}$$

$$\Phi(moveCPL_{igp}) = \begin{cases} -300, & \text{if} \quad moveCPL_{igp} < 300, \\ moveCPL_{igp}, & \text{if} \quad moveCPL_{igp} \in [-300,300], \\ 300, & \text{if} \quad moveCPL_{igp} > 300. \end{cases}$$

The binary variable d_{igp} equals one if both, the move played and the move suggested by the engine, are evaluated within the interval $[-200, 200]$. Only in this case will the evaluation be considered for the game-level skill measure. The reason for this is that players who have a strong advantage might not choose to play the best but a less risky move and therefore would get punished illegitimately. If a player considers his position to be lost, he similarly might on purpose play a worse move in an attempt to turn around the advantage, which again would result in an objectively illegitimate evaluation loss.

Additionally, for some mistakes a player makes, he gets punished unproportionately high, resulting from the numerical difference of his move to the evaluated choice of the engine. Guid and Brakto (2011) therefore confine the output space of the evaluation function to the interval $[-300, 300]$, where 300 CPLs should represent a "huge mistake [...] through the eyes of a human expert."