Average Controllability Measures for Solitaire Games*

Marco Faella
Dept. of Electrical Engineering and Information Technologies
University of Naples "Federico II", Italy
m.faella@unina.it

ABSTRACT

We discuss several ways to measure the extent to which a player can exert control over a one-player game, relating them to the existing literature on the "skill vs chance" dichotomy. We focus on measures that depend only on the rules of the games, and not on how people actually play them. After presenting a set of desirable properties, we show that two statistical measures of effect size satisfy them and we estimate the value of such measures on several well-known games.

Keywords

Markov decision processes, solitaire games, skill vs luck

1. INTRODUCTION

Many games played by people feature a combination of determinism and chance. Most card games involve an initial shuffling, board games resort to dice, and electronic games employ pseudo-random number generators.

There is no established way to measure the amount of chance involved in a game, besides some isolated efforts that will be discussed in Section 4. By measuring the amount of chance, we mean assign a real number between 0 and 1 to every game, in such a way that a purely deterministic game, such as Chess, receives value 1, and a purely random game, such as a lottery, receives value 0. In the spirit of our discipline, we are interested in measures that can be automatically computed, or at least approximated to an arbitrary precision. Hence, the candidate measures may depend only on the rules of the game, and not on how people actually play it. This property has the additional advantage that the measures we consider can also be evaluated on games that are not (yet) played by people, perhaps because they are still being designed.

The motivation that drives most of the literature on this subject comes from gambling laws. Many countries distinguish games of chance from games of skill and apply different regulations to each class. Such laws make no attempt to formally define these two classes, leaving this task to judges, who may very well hold different opinions on the subject. In fact, in at least one well-documented case, different courts have expressed contrasting rulings on the game of electronic draw poker (a.k.a. video poker) [13], which is one of the games we examine in Section 7. Our investigation may clearly contribute to these applications, by providing an objective and *a-priori* means to quantify the amount of chance in a game.

Measures of chance can also help players make informed decisions on which games to play according to their tastes, and enable game designers to tune their games for their audience and facilitate innovative game designs. In the words of game designer Soren Johnson [12]: "The appropriate role of chance in a game is ultimately a subjective question, and giving players the ability to adjust the knobs [...] can open up the game to a larger audience with a greater variety of tastes."

Last but not least, we believe that the problem is interesting for its own sake, because it bears connections with different disciplines, ranging from learning theory to psychology. In this paper, we focus on *solitaire games*. Common solitaire games include video poker, Black Jack [1], Minesweeper, etc. Such games can be modeled as finite-state discrete-time Markov Decision Processes (MDPs). We can then rephrase our original problem as a discussion of measures of *controllability* for MDPs. To the best of our knowledge, the only related work focused on MDPs is [21]. There, the authors look for metrics on MDPs that may help choose the best learning algorithm for a planning problem. They propose several metrics and experimentally evaluate their ability to predict the performance of different learning algorithms. The intuition behind their proposals is similar to ours: measure the degree of controllability of the model. However, the purpose and methodology is very different, leading to a number of metrics that have a "local" flavor, in the sense that they are easy to compute based on the information which is available in a state. It is possible that our investigation may prove useful in that domain, although the metrics we propose are likely to be much more computationally demanding than theirs.

An entirely different domain where similar concerns are routinely addressed is in *effect-size statistics*. In many social and life sciences, investigators try to measure the influence of some controllable (or at least measurable) variable on the outcome of a complex process [18]. This is the familiar

*This work was supported by the European Union 7th Framework Program project 600958-SHERPA.

Appears in: *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2016), J. Thangarajah, K. Tuyls, C. Jonker, S. Marsella (eds.), May 9–13, 2016, Singapore.*

[1]In Black Jack, the dealer follows a deterministic strategy, so a game between a single player and the dealer is effectively a solitaire.

scenario in which a control group is compared with one or more treatment groups, to evaluate the effect of a drug, a procedure, or some intrinsic attribute of the subjects. The same methodology also applies to Computer Science, when one wants to accurately characterize the influence of some feature of the input on the performance of an algorithm [3]. In most cases, the complex process which transforms inputs into an outcome is treated as a black-box, and statistics is used to analyze the relation between inputs and outputs.

Analogously, the outcome of a solitaire game is influenced by some controllable input (the player's moves) and a random device. Contrary to the complex systems of Biology, Medicine, and Psychology, the process that combines controllable inputs and random events is exactly known, being determined by the rules of the game. On the other hand, the state-space of a real game can be so large that statistical techniques remain the only applicable analysis tool. Effect-size statistics constitute the standard toolbox for analyzing these scenarios. Many different effect-size measures can be found in the literature [9], some of which will be presented in Section 5, and used in Section 6 as measures of average controllability for solitaire games.

The reason for targeting *average* controllability, rather than minimum or maximum controllability, stems from the observation that the first may more accurately approximate the behavior of a population of players of various skill levels. After all, min, average, and max are the three canonical points in the spectrum of rationality, with bounded rationality and actual human players arguably positioned somewhere between the second and the third point.

Summarizing, after the preliminary notions of Section 2, **our contribution** is organized in four sections, followed by some conclusions. In Section 3, we present a small set of axioms that we deem desirable for a measure of controllability. In fact, the property we are trying to measure appears so nebulous that little to no consensus has been reached by the scientific community on its interpretation. Hence, we believe that it is useful to start by enumerating some axioms that focus our investigation and ease the comparison with the related literature. Such comparison is indeed the topic of Section 4, which reviews previous proposals of controllability measures, putting them in the context of solitaire games and analysing them from a computational point of view. We conclude that previous measures are not suitable to our aims, and therefore devote Sections 5 and 6 to outline our proposal, which is based on the effect-size measures "percentage of variance explained" and "percentage of absolute deviation explained". We show that such measures satisfy a basic set of axioms but violate a stronger version of one of the axioms. Then, we prove that these measures define two different orders between games.

In Section 7, we employ a purpose-built game simulation framework to estimate the value of each measure on randomized versions of Tic-tac-toe and simplified variants of Minesweeper, Video Poker, and Black Jack. Moreover, we include in the comparison an individual sport, archery, for which we obtained experimental data from an external source.

2. DEFINITIONS

We represent a solitaire game by a finite-state discrete-time MDP with partial information and a reachability objective.

Definition 1. A *Markov decision process with partial information (PIMDP)* is a tuple $P = (S, s_0, I, A, \Gamma, \delta)$, where S is a finite set of states, $s_0 \in S$ is the initial state, I is a partition of S into information sets, A is a finite set of actions, $\Gamma : I \to 2^A$ gives the set of active actions for each information set, and $\delta : S \times A \times S \to [0, 1]$ is the transition function, assigning to each triple (s, a, s') the probability that the process moves from s to s' when action a is selected. Denote by $[s]$ the information set containing s, for all $s \in S$ and $a \in \Gamma([s])$, it holds $\sum_{s' \in S} \delta(s, a, s') = 1$.

PIMDPs are a special case of partially observable MDPs [7] and are similar to the game theoretic notion of extensive game with imperfect information [19]. In particular, states belonging to the same information set are indistinguishable to the player, so that he is forced to take the same action in all of them.

A *path* in P is a finite sequence of states $s_0 s_1 \dots s_n$, such that for all $i = 0, 1, \dots, n-1$ there exists $a \in \Gamma([s_i])$ such that $\delta(s_i, a, s_{i+1}) > 0$. A path is *maximal* if it is not a proper prefix of another path (because its last state has no active actions). For our purposes, it suffices to consider *acyclic* PIMDPs, i.e., PIMDPs having no path with repeated states. Notice that, in order for the process to be acyclic, some states must have no active actions (sink states).

A (memoryless and deterministic) *strategy* is a function $\sigma : I \to A$ such that $\sigma(\mathcal{I}) \in \Gamma(\mathcal{I})$ for all $\mathcal{I} \in I$. A path $s_0 s_1 \dots s_n$ is *consistent* with the strategy σ if for all $i = 0, 1, \dots, n$, it holds that $\delta(s_i, \sigma(s_i), s_{i+1}) > 0$. We denote by $runs(\sigma)$ the set of all maximal paths consistent with σ. As we are assuming that the process is acyclic, each strategy induces a finite set of maximal paths and a discrete probability space over it in the usual way (see [7] for details). We denote by p_σ the probability measure induced by σ. If two strategies give rise to the same probability space, we consider them equivalent. So, when talking about the set of all strategies in a PIMDP, we implicitly refer to the quotient w.r.t. such equivalence.

Definition 2. A *solitaire game* G is a pair (P, T) where P is an acyclic PIMDP and $T \subseteq S$ is a non-empty set of target states.

Given a solitaire game (P, T), assume w.l.o.g. that states in T are sinks, and let *val* be the function assigning 1 to the paths ending in T, and 0 to all other paths. We denote by E_σ the probability of reaching T under strategy σ. Notice that E_σ is also the expected value of the random variable assigning value $val(r)$ to each run $r \in runs(\sigma)$, under the probability space induced by σ. So, we also call it the *expected value* of σ.

3. DESIRED QUALITIES OF A MEASURE

In this section, we discuss a number of desired properties for a controllability measure for solitaire games, having in mind the applications mentioned in the Introduction. First of all, we choose to restrict to games that can be put in the form of a PIMDP.

Axiom 0. The measure can only depend on the rules of the game, as formalized by Definition 2.

Although Axiom 0 sounds perfectly natural in a computational context, it sets us apart from the previous research on the topic (see Section 4), so let us remark once more its benefits. Axiom 0 makes sure that the measure can be applied to any game that can be put in the form of Definition 2, allowing us to evaluate hypothetical games, or a proposed variant to a real-world game, and ascertain whether it would increase or decrease the randomness of the game.

We say that a game is *chance-less* if there is a single destination for each action, i.e., for all $s \in S$ and $a \in \Gamma([s])$ there exists a unique s' such that $\delta(s, a, s') > 0$. Dually, a game is *skill-less* if there is a single active action from each information set: for all $\mathcal{I} \in I$ it holds that $|\Gamma(\mathcal{I})| = 1$. A skill-less game has a single strategy. Intuitively, all chance-less (resp., skill-less) games should have controllability measure 1 (resp., 0). However, some degenerate games are chance-less and skill-less at the same time, giving rise to a single maximal path. What should be their measure? We stipulate that it should be 0, as our primary objective is to measure to what extent the moves of the player decrease the uncertainty in the outcome. If a game has no uncertainty to begin with, no control can be exerted by the player.

Axiom 1. For all chance-less games G, we have $m(G) = 1$, provided that there are at least two runs r_1, r_2 in the game, with $val(r_1) \neq val(r_2)$.

Axiom 2. For all skill-less games G, we have $m(G) = 0$.

Axioms 1 and 2 are entirely uncontroversial and can be found in all previous works on the subject.

Next, we introduce an axiom which provides a canonical game for each possible value of the measure. The family of games $B_p(G_0, G_1)$ is depicted in Figure 1(c) and accepts as parameters a probability $p \in [0, 1]$ and two other games G_0 and G_1. Games *Rand* and *Ctrl* are shown in Figures 1(b) and 1(a), respectively.

Axiom 3. For all $p \in [0, 1]$, we have $m\big(B_p(\mathit{Rand}, \mathit{Ctrl})\big) = p$.

Axiom 3 fixes the value of the measure for the family of games that are simple mixtures of a specific game of measure 0 and a specific game of measure 1. It is tempting to replace it with one of the following increasingly stronger versions. In the first version, games *Rand* and *Ctrl* are replaced by arbitrary games having measure 0 and 1, respectively.

Axiom 3'. For all $p \in [0, 1]$, let G_0 and G_1 be arbitrary games such that $m(G_0) = 0$ and $m(G_1) = 1$. Then, we have $m\big(B_p(G_0, G_1)\big) = p$.

Notice that Axiom 3' introduces a discontinuity in the measure, in the sense that we can exhibit a sequence of games having measure $\frac{1}{2}$, converging to a game that has measure 1 due to Axiom 1. For $i > 0$, the i-th game in the sequence is $B_{1/2}(H_i, \mathit{Ctrl})$, where H_i is a skill-less game in which the target T is reached with probability $\frac{1}{2^i}$. In itself, this discontinuity should not be a reason to discard Axiom 3'. After all, even in a simple Markov chain (with loops), the probability of eventually reaching a given state is discontinuous w.r.t. the one-step transition probabilities. In the next section, we show that two classical effect size measures recalled in Section 5 satisfy Axiom 3 but not Axiom 3'.

In another, even stronger version of Axiom 3, games G_0 and G_1 are entirely arbitrary.

Axiom 3". For all $p \in [0, 1]$, let G_0 and G_1 be arbitrary games. Then, we have

$$m\big(B_p(G_0, G_1)\big) = p \cdot m(G_1) + (1 - p)m(G_0).$$

We now show that Axiom 3" is too strong for our intents. Consider the game $B_{1/2}(G_0, G_1)$ where $G_0 = G_1 = \mathit{Ctrl}$. After an initial random move, the player is in total control and can choose between winning and losing. Axiom 3" prescribes measure 1 for this game, because the sub-games G_0 and G_1 are entirely controlled by the player and hence have measure 1 due to Axiom 1. On the other hand, if we identify (types of) players with strategies, this game has four different strategies, depending on the choice of the player in the two sub-games. If the player chooses to win when in G_0 and lose in G_1 (or vice versa), then the outcome is entirely decided by the initial random move. In other words, 2 out of 4 strategies effectively exert no control on the game. For this reason, we reject Axiom 3".

The decision to consider all possible strategies in the game stems from our objective of defining *average* controllability measures. In so doing, we violate the classical principle of Independence of Irrelevant Alternatives. Indeed, we assume that players, due to their bounded rationality, may fail to recognize that the two sub-games G_0 and G_1 are identical, and hence may behave differently in the two sub-games. A similar tension between bounded rationality and the classic tenets of decision theory is noted by Dagsvik [4].

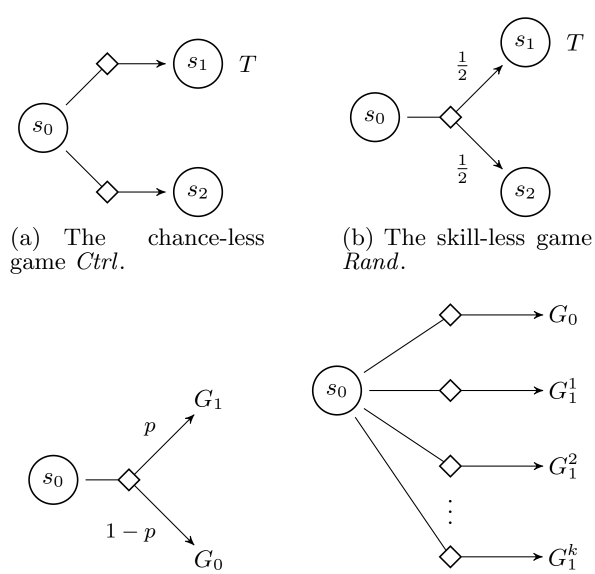

(a) The chance-less game *Ctrl*.

(b) The skill-less game *Rand*.

(c) The game family $B_p(G_0, G_1)$, with $p \in [0, 1]$.

(d) The game family $A_k(G_0, G_1^1, \ldots, G_1^k)$, with $k \in \mathbb{N}$.

Figure 1: Significant games. States are circles and actions are diamonds.

A simple measure that satisfies Axioms 0-3 as well as Axiom 3' is the difference between the expected value of the optimal strategy and the expected value of the worst strategy. However, we believe that this measure is too crude to be useful to our inteded applications. In particular, consider two games G', G'' with a thousand strategies each. In G', 999 strategies have expected value $1/2$, which denotes no control on the game outcome, and one strategy has expected value 1. In G'', the opposite happens: 999 strategies win with certainty and one strategy has expected value $^1/_2$. The measure described above assigns the same measure $^1/_2$ to both games, whereas we would like $m(G'') > m(G')$. We therefore introduce the following axiom, based on the family of games A_k in Figure 1(d):

Axiom 4. For all $k > j > 0$, we have $m\big(A_k(\mathit{Rand}, \mathit{Ctrl}, \ldots, \mathit{Ctrl})\big) > m\big(A_j(\mathit{Rand}, \mathit{Ctrl}, \ldots, \mathit{Ctrl})\big)$.

Axiom 4 states that adding more alternatives for the player that lead to completely controllable games raises the value of the measure (even if other controllable alternatives are already available). Essentially, Axiom 4 implies that we are looking for an *average* measure of controllability, rather than the maximum amount of control that a player can exercise. Similarly to Axiom 3, we also consider a stronger version of Axiom 4, which employs generic games with prescribed measures instead of the two canonical games *Rand* and *Ctrl*.

Axiom 4'. For all $k > j > 0$, let $G_0, G_1^1, \ldots, G_1^k$ be games such that $m(G_0) = 0$, there is at least one strategy σ in G_0 s.t. $E_\sigma \in (0,1)$, and $m(G_1^i) = 1$ for all $i = 1, 2, \ldots, k$. Then, we have $m\big(A_k(G_0, G_1^1, \ldots, G_1^k)\big) > m\big(A_j(G_0, G_1^1, \ldots, G_1^j)\big)$.

The second constraint on G_0 ensures that there is some uncertainty if the player chooses that game.

In the next section, we review, in the light of these axioms, the main proposals that have been made in the literature.

4. PREVIOUS APPROACHES

In the literature, the issues we are considering here are often discussed under the "skill vs luck" dichotomy [17]. We believe that such terminology is somewhat misleading, because it may contribute to confounding two different and largely independent issues: the amount of randomness in the rules of a game (the "chance" element) and the difficulty of a game for a human player (the "skill" element). Consider Tic-Tac-Toe and Chess: the rules of both games involve no randomness, but nobody would say that they require the same amount of skill to be played well. Our Axiom 1, shared by similar works in the literature, makes it clear that we are interested in the amount of randomness and not in the difficulty of winning the game. Which (combination) of these two issues should be of interest to the legislator is up for debate and we are in no position to contribute. Here, we focus on measuring the amount of randomness in the rules of a game, and we call *controllability* the lack thereof.

Borm and van der Genugten [2], with later contributions by Dreef [5, 6], seem to be the first to formally define a measure of controllability for games, coming from an Economic background, with gambling laws as the reference application. They subscribe to our Axioms 1 and 2 and propose the following measure:

$$m_1(G) = \frac{\text{learning effect}}{\text{learning effect} + \text{random effect}}.$$

Roughly speaking, the learning effect quantifies the difference between the performance of an expert player and that of a beginner, while the random effect quantifies the advantage given by knowing in advance the realization of the random effects in the game.

The authors then clarify the above definition by semi-formally defining the following three types of players: *(i)* the *optimal* player, who plays the strategy with the highest expected payoff, or, in our setting, the highest probability of reaching the target T; *(ii)* the *fictive* player, who plays optimally and moreover knows in advance the realization of the random events that will occur during play; *(iii)* the *beginner* player, who "has only just familiarized himself with the rules of the game" [5].

Call E_{opt}, E_{fic}, and E_{beg} the expected payoff of the three player categories outlined above, Dreef et al. define the learning effect as $E_{\text{opt}} - E_{\text{beg}}$ and the random effect as $E_{\text{fic}} - E_{\text{opt}}$, leading to the following expression for m_1:

$$m_1(G) = \frac{E_{\text{opt}} - E_{\text{beg}}}{(E_{\text{opt}} - E_{\text{beg}}) + (E_{\text{fic}} - E_{\text{opt}})} = \frac{E_{\text{opt}} - E_{\text{beg}}}{E_{\text{fic}} - E_{\text{beg}}}.$$

Even in the context of the original paper, the authors struggle to define the meaning of the beginner player, discussing three possible ways to ascertain his behavior: *(a)* by postulating that she chooses uniformly at random among the available moves; *(b)* by statistically analyzing the actual behavior of people playing the game; *(c)* by assigning the task to a "gambling expert". In their analysis of *minipoker*, Dreef et al. employ option *(c)*, i.e., they devise an ad-hoc strategy that is deemed appropriate for a beginner. On the other hand, only option (a) satisfies our Axiom 0 and is suitable for a fair comparison between different games.

Regardless of the interpretation chosen for the beginner player, the biggest problem with measure m_1 is that it violates Axiom 3. Indeed, in the game family B_p, knowing in advance the result of the random moves gives *no* advantage to the player. Precisely, for game $B_p(\mathit{Rand}, \mathit{Ctrl})$ we have that $E_{\text{opt}} = E_{\text{fic}} = \frac{p+1}{2}$, and therefore $m_1\big(B_p(\mathit{Rand}, \mathit{Ctrl})\big) = 1$ for all $p \in [0,1]$. This critique to measure m_1 was already raised by Heubeck [10].

On the other hand, if we adopt interpretation (a) for "beginner" (i.e., "uniformly random"), the simple expression $m_2(G) = 2(E_{\text{opt}} - E_{\text{beg}})$, which is twice the "learning effect", gives the desired value p for the game $B_p(\mathit{Rand}, \mathit{Ctrl})$, thus satisfying Axiom 3. However, it is easy to see that m_2 violates Axiom 1. Indeed, let G be a chance-less game with three strategies, two of which lead to the target T. We have $m_2(G) = {}^2/_3$ rather than the prescribed value of 1.

It should be noted that Dreef et al. [5] mention in their conclusions that an alternative approach could be based on the analysis of variance, which is one of the central tenets of the present paper.

In a recent column [22], Wyner presents a critique of Deer et al. and advances his own proposal, accompanied by a set of (semi-formal) axioms he deems desirable. The first two axioms unsurprisingly coincide with our Axioms 1 and 2. We quote the remaining four axioms, slightly adapting them to our notation:

Axiom III. If G^N is the convolution of N rounds of game G, then $\lim_{N\to\infty} m(G^N) = 1$ provided that $m(G) > 0$.

The paper does not define "convolution", but we can surmise it means playing the game multiple times and somehow combine the obtained payoffs. The intended meaning of the axiom is that if a game involves a non-zero amount of skill, in the long-run its outcome will depend entirely on skill. The axiom makes sense for 2-player games: in the long run, chance elements should even out and the different skill (or, in more neutral terms, different strategies) of the two players is what really counts towards establishing a winner. In fact, a similar principle was used in recent work on online multi-player poker [20].

However, applying this axiom to controllability of solitaires appears problematic. In a game with non-zero controllability, the moves of the player have *some* influence on the distribution of outcomes. On the other hand, having perfect control (measure 1) means that there is no randomness in the game and the game outcome is directly determined by the player's moves. If a game with non-zero controllability measure is repeated at will, there is no reason to believe that the outcome has no randomness left.

Axiom IV. The measure should depend on the distribution of skill levels of the population of players.

Although perhaps motivated in a context which is strictly focused on legal applications, this axiom goes squarely against our Axiom 0 and the idea of achieving a fair comparison between different games.

The following axiom is a weaker form of our Axiom 3. In fact, the game family B_p from Figure 1(c) is a straightforward generalization of the *balanced mixture game* of Wyner, which coincides with $B_{1/2}$.

Axiom V. It holds $m\big(B_{1/2}(Rand, Ctrl)\big) = \frac{1}{2}$.

The final axiom refers to the fact that several juristictions define games of skill as those games where skill plays a preponderant role, leading to the so-called *preponderance test*. Axiom VI suggests to set the threshold at $1/2$. We briefly comment on this axiom in Section 7, after we present our experimental evaluations.

Axiom VI. Any game G for which $m(G) < \frac{1}{2}$ fails the preponderance test.

After presenting these axioms, Wyner introduces his proposal for a measure of skill in multi-player games. Roughly speaking, he postulates the possibility to quantify the skill level of a player. Then, for two skill levels $k > j$, a match between two players is termed (j, k)*-reversible* if its outcome is reversed when the skill level of the losing player changes from j to k (and the random events somehow remain fixed). The proposed measure is the probability that a match is reversible "when averaged over players of random talent against experts" [22]. Clearly, such semi-formal definition is hard to fit into the current framework, particularly due to our Axiom 0 and the presence of a single player.

Heubeck [11] discusses measures similar to the percentage of absolute deviation explained for games of a special form, in which a completely controllable phase is followed by a completely random choice of payoff.

Finally, a significant amount of attention has been devoted to the balance between skill and chance in (multi-player) Poker [16, 20]. Exploiting the wealth of data currently available on online poker, these works take an empirical approach and estimate the correlation between past and future performance of human players.

5. EFFECT-SIZE MEASURES

We recall the basic definitions pertaining the family of effect-size measures called "percentage of variance explained", commonly associated to the statistical framework ANOVA (analysis of variance) [8]. These measures aim at attributing the spread in the outcomes to two different sources: the strategy giving rise to each outcome and the variation within a single strategy. To compute such a measure, one summarizes each strategy σ by its expected value E_σ; then, one quantifies the spread of such averages and divides it by the total spread, where "spread" can be interpreted in different ways.

Given a set of strategies $\sigma_1, \dots, \sigma_n$, let $\overline{E}$ be the arithmetic mean of $E_{\sigma_1}, \dots, E_{\sigma_n}$, i.e., $\overline{E} = \frac{1}{n}\sum_{i=1}^n E_{\sigma_i}$. We obtain the first measure by interpreting "spread" as variance, leading to the notions of *variance between strategies* Var_{btw} and *total variance* Var_{tot}:

$$Var_{\mathrm{btw}} = \frac{1}{n}\sum_{i=1}^n (E_{\sigma_i} - \overline{E})^2$$

$$Var_{\mathrm{tot}} = \frac{1}{n}\sum_{i=1}^n \sum_{r \in runs(\sigma_i)} (val(r) - \overline{E})^2 \cdot p_{\sigma_i}(r).$$

The classical effect size measure η^2, called *percentage of variance explained* is the ratio of the above two statistics:

$$\begin{aligned}
\eta^2 = \frac{Var_{\mathrm{btw}}}{Var_{\mathrm{tot}}} &= \frac{\sum_{i=1}^n (E_{\sigma_i} - \overline{E})^2}{\sum_{i=1}^n \sum_{r \in runs(\sigma_i)} (val(r) - \overline{E})^2 \cdot p_{\sigma_i}(r)} \\
&= \frac{\sum_{i=1}^n \left(E_{\sigma_i}^2 + \overline{E}^2 - 2\overline{E}E_{\sigma_i}\right)}{\sum_{i=1}^n \left(E_{\sigma_i} + \overline{E}^2 - 2\overline{E}E_{\sigma_i}\right)} \\
&= \frac{\sum_{i=1}^n E_{\sigma_i}^2 + n\overline{E}^2 - 2\overline{E}\sum_{i=1}^n E_{\sigma_i}}{\sum_{i=1}^n E_{\sigma_i} + n\overline{E}^2 - 2\overline{E}\sum_{i=1}^n E_{\sigma_i}} \\
&= \frac{\sum_{i=1}^n E_{\sigma_i}^2 - n\overline{E}^2}{n(\overline{E} - \overline{E}^2)}.
\end{aligned}$$

The value of η^2 is undefined when both its numerator and its denominator are zero, which in our setting happens if and only if $\overline{E} \in \{0, 1\}$. Since this corresponds to games whose outcome is independent from the player strategy, in this case we set the value of η^2 to 0. We will also be interested in the square root of η^2, naturally denoted by η, which can be interpreted as the percentage of standard deviation which is explained by the choice of a strategy: $\eta = \sqrt{\eta^2} = \frac{\sqrt{Var_{\mathrm{btw}}}}{\sqrt{Var_{\mathrm{tot}}}}$.

If instead "spread" is interpreted as absolute deviation, we obtain the following definitions of *absolute deviation between strategies* and *total absolute deviation*.

$$AD_{\mathrm{btw}} = \frac{1}{n}\sum_{i=1}^n |E_{\sigma_i} - \overline{E}|$$

$$AD_{\mathrm{tot}} = \frac{1}{n}\sum_{i=1}^n \sum_{r \in runs(\sigma_i)} |val(r) - \overline{E}| \cdot p_{\sigma_i}(r).$$

Based on the above, the *percentage of absolute deviation explained* γ can be defined as follows:

$$\begin{aligned}
\gamma = \frac{AD_{\mathrm{btw}}}{AD_{\mathrm{tot}}} &= \frac{\sum_{i=1}^n |E_{\sigma_i} - \overline{E}|}{\sum_{i=1}^n \sum_{r \in runs(\sigma_i)} |val(r) - \overline{E}| \cdot p_{\sigma_i}(r)} \\
&= \frac{\sum_{i=1}^n |E_{\sigma_i} - \overline{E}|}{2n\big(\overline{E} - \overline{E}^2\big)}.
\end{aligned}$$

As before, we set $\gamma = 0$ whenever $AD_{\mathrm{btw}} = AD_{\mathrm{tot}} = 0$, which happens iff $\overline{E} \in \{0, 1\}$.

6. CONTROLLABILITY MEASURES

In this section, we show that the two effect size measures η and γ satisfy the basic axioms of Section 3, but fail to satisfy the stronger Axiom 3'.

By an abuse of notation, for a game G we write $\eta(G)$ and $\gamma(G)$ to signify $\eta(E_{\sigma_1},\dots,E_{\sigma_n})$ and $\gamma(E_{\sigma_1},\dots,E_{\sigma_n})$, where $\sigma_1,\dots,\sigma_n$ are all the strategies in G, modulo equivalence.

First, we characterize the cases in which the two measures have value 1.

PROPOSITION 1. *Let $E_1,\dots,E_n \in [0,1]$, and $\overline{E}$ be their arithmetic average. The following are equivalent:*

1. *$\eta(E_1,\dots,E_n) = 1$;*
2. *$\gamma(E_1,\dots,E_n) = 1$;*
3. *$E_i \in \{0,1\}$ for all $i = 1,\dots,n$, and $\overline{E} \in (0,1)$.*

PROOF . [(1)↔(3)]. It holds that

$$\begin{aligned}\eta = 1 &\Leftrightarrow \textstyle\sum_i E_i^2 - n\overline{E}^2 = n\overline{E} - n\overline{E}^2\\ &\Leftrightarrow \textstyle\sum_i E_i^2 = n\overline{E}\\ &\Leftrightarrow \textstyle\sum_i E_i^2 = \sum_i E_i. \qquad (*)\end{aligned}$$

As $E_i \in [0,1]$, it follows that $E_i^2 \le E_i$. By (*), we have that $E_i^2 = E_i$ for all i. Since $E_i^2 = E_i$ iff $E_i \in \{0,1\}$, we obtain the thesis.

[(2)↔(3)]. Let $S = \sum_i E_i = n\overline{E}$ and let I^+ (resp., I^-) be the set of indices $i \in \{1,\dots,n\}$ such that $E_i \ge \overline{E}$ (resp., $E_i < \overline{E}$). Assuming that $\gamma = 1$, by definition we have that $\overline{E} \in (0,1)$. Moreover,

$$\begin{aligned}\textstyle\sum_{i\in I^+}(E_i - \overline{E}) + \sum_{i\in I^-}(\overline{E} - E_i) &= 2n\overline{E}(1-\overline{E});\\ \sum_{i\in I^+} E_i - |I^+|\cdot\overline{E} + |I^-|\cdot\overline{E} - \sum_{i\in I^-} E_i &= 2S(1-\overline{E});\\ \frac{\sum_{i\in I^+} E_i}{S} - \frac{\sum_{i\in I^-} E_i}{S} + \frac{|I^-|}{n} - \frac{|I^+|}{n} &= 2(1-\overline{E});\\ 1 - 2\frac{\sum_{i\in I^-} E_i}{S} + 1 - 2\frac{|I^+|}{n} &= 2 - 2\overline{E};\\ \frac{|I^+|}{n} + \frac{\sum_{i\in I^-} E_i}{S} &= \overline{E};\\ |I^+| + \frac{\sum_{i\in I^-} E_i}{\overline{E}} &= \sum_{i\in I^+} E_i + \sum_{i\in I^-} E_i.\end{aligned}$$

Now, write the last equation as $a + b = c + d$, it is easy to see that $a \ge c$ and $b \ge d$. Hence, it holds $a = c$ and $b = d$, which implies (3).

Conversely, assume that (3) holds. We have that $\overline{E} = \frac{|I^+|}{n}$ and hence

$$\begin{aligned}\gamma &= \frac{\sum_{i\in I^+}(1-\overline{E}) + \sum_{i\in I^-}\overline{E}}{2n(\overline{E} - \overline{E}^2)}\\ &= \frac{|I^+| - |I^+|\cdot\overline{E} + |I^-|\cdot\overline{E}}{2n(\overline{E} - \overline{E}^2)}\\ &= \frac{|I^+| - |I^+|\cdot\frac{|I^+|}{n} + (n - |I^+|)\cdot\frac{|I^+|}{n}}{2n(\frac{|I^+|}{n} - \frac{|I^+|^2}{n^2})}\\ &= \frac{2|I^+| - 2\frac{|I^+|^2}{n}}{2|I^+| - 2\frac{|I^+|^2}{n}} = 1.\end{aligned}$$

□

PROPOSITION 2. *Measure η satisfies Axioms 0, 1, 2, 3, and 4'.*

PROOF. Axiom 1 is an immediate consequence of Prop. 1 whereas Axiom 2 follows from the definition of η. Axiom 3 can be verified by simple calculations. As for Axiom 4', let $G_0, G_1^1, G_1^2, \dots, G_1^k$ be games such that $\eta(G_0) = 0$ and $\eta(G_1^i) = 1$ for all $i = 1,\dots,k$. We prove that $\eta\big(A_k(G_0, G_1^1,\dots,G_1^k)\big) > \eta\big(A_{k-1}(G_0, G_1^1,\dots,G_1^{k-1})\big)$. By Prop. 1, all strategies in G_1^k have expected value 0 or 1. Hence, it is sufficient to prove that for all $E_1,\dots,E_n \in [0,1]$, *(i)* $\eta(E_1,\dots,E_n,0) > \eta(E_1,\dots,E_n)$ and *(ii)* $\eta(E_1,\dots,E_n,1) > \eta(E_1,\dots,E_n)$. The thesis then follows by an induction on the number of strategies in the sub-game G_1^k.

We develop the proof for (i), as the one for (ii) is similar. For technical convenience, we prove the result for η^2, as the thesis for η follows as a consequence. Let $\overline{E}$ (resp., S) be the arithmetic average (resp., the sum) of $(E_1,\dots,E_n)$ and let $\overline{E}'$ be the arithmetic average of $(E_1,\dots,E_n,0)$.

$$\begin{aligned}&\eta^2(E_1,\dots,E_n,0) - \eta^2(E_1,\dots,E_n) =\\ &\quad\frac{\sum_i E_i^2 - (n+1)\overline{E}'^2}{(n+1)\overline{E}'(1-\overline{E}')} - \frac{\sum_i E_i^2 - n\overline{E}^2}{n\overline{E}(1-\overline{E})} =\\ &\quad\frac{(1-\overline{E})(\sum_i E_i^2 - \frac{S^2}{n+1}) - (1-\overline{E}')(\sum_i E_i^2 - \frac{S^2}{n})}{S(1-\overline{E}')(1-\overline{E})}.\end{aligned}$$

Since the denominator is clearly non-negative, we can focus on the sign of the numerator.

$$\begin{aligned}&(1-\overline{E})(\textstyle\sum_i E_i^2 - \dfrac{S^2}{n+1}) - (1-\overline{E}')(\textstyle\sum_i E_i^2 - \dfrac{S^2}{n}) =\\ &\textstyle\sum_i E_i^2 - \dfrac{S^2}{n+1} - \overline{E}\textstyle\sum_i E_i^2 + \dfrac{S^3}{n(n+1)}\\ &\qquad - \textstyle\sum_i E_i^2 + \dfrac{S^2}{n} + \overline{E}'\textstyle\sum_i E_i^2 - \dfrac{S^3}{n(n+1)} =\\ &S^2\Big(\frac{1}{n} - \frac{1}{n+1}\Big) - S\textstyle\sum_i E_i^2\Big(\dfrac{1}{n} - \dfrac{1}{n+1}\Big) =\\ &S(S - \textstyle\sum_i E_i^2)\Big(\dfrac{1}{n} - \dfrac{1}{n+1}\Big).\end{aligned}$$

Under the assumptions of Axiom 4', all three terms of the last expression are strictly positive, which proves our thesis. □

PROPOSITION 3. *Measure γ satisfies Axioms 0, 1, 2, 3, and 4.*

PROOF. Axiom 1 is an immediate consequence of Prop. 1 whereas Axiom 2 follows from the definition of γ. Axioms 3 and 4 can be verified by simple calculations. In particular, $\gamma\big(A_k(\mathit{Rand}, \mathit{Ctrl},\dots,\mathit{Ctrl})\big) = \frac{2k}{k+1}$, which is strictly increasing in k, for $k \ge 1$. □

PROPOSITION 4. *Measures η and γ violate Axiom 3'.*

PROOF. Consider the game $G_\varepsilon = B_{1/2}(H_\varepsilon, \mathit{Ctrl})$, where H_ε is a skill-less game whose only strategy has expected value $\varepsilon \in [0,1]$. Axiom 3' prescribes the measure of G to be $\frac{1}{2}$, for all $\varepsilon \in [0,1]$. On the other hand, simple calculations show that $\eta(G_\varepsilon) = (3 + 4\varepsilon - 4\varepsilon^2)^{-\frac{1}{2}}$ and $\gamma(G_\varepsilon) = (\frac{3}{2} + 2\varepsilon - 2\varepsilon^2)^{-1}$, which are both different from $\frac{1}{2}$ when $\varepsilon \neq \frac{1}{2}$. □

Since our measures span a conventional range from 0 to 1, we are not primarily interested in the actual value assigned to a given game, but rather in the weak order induced on games by the measure. A measure m induces the order $\leq_m$ in the following way: $G_1 \leq_m G_2$ iff $m(G_1) \leq m(G_2)$. Such order is weak because there may be games with the same measure, that end up being equivalent for the corresponding order. Two measures are called *co-monotonic* if they induce the same order. A simple example shows that the two measures we are considering are *not* co-monotonic.

PROPOSITION 5. *Measures η and γ are not co-monotonic.*

PROOF. Consider two games G_1, G_2, with 3 strategies each. The expected values of the strategies and the approximate value of all measures are reported in the following table.

game	E_a	E_b	E_c	η	γ
G_1	0	0.2	0.4	0.41	0.417
G_2	0	0.31	0.31	0.36	0.420

Moving from G_1 to G_2, γ increases while η decreases, thus proving the thesis. $\square$

The next section contains more extreme examples of non-co-monotonicity, such as those arising in the analysis of the game Minesweeper.

7. EXPERIMENTAL EVALUATION

To develop an intuitive understanding of our measures, we set up a game simulation framework and implemented various well-known games in it. The main results are presented in Figure 2.

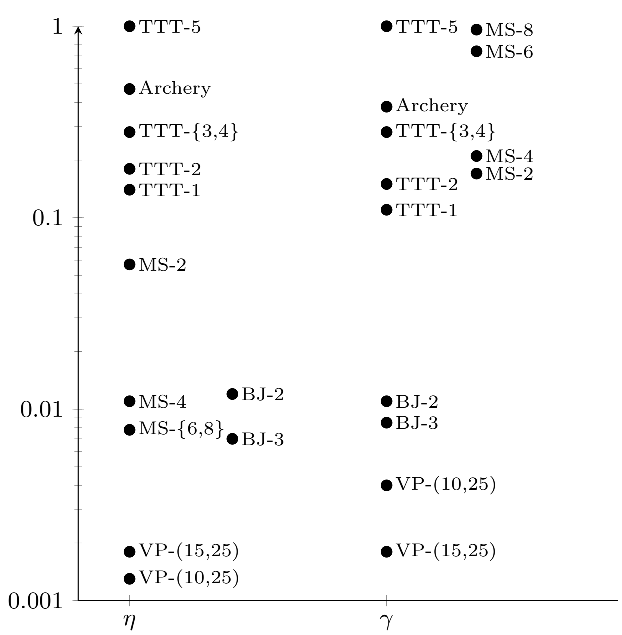

Figure 2: The amount of controllability in games Tic-tac-toe (TTT), Black Jack (BJ), Video Poker (VP), Minesweeper (MS), and an instance of archery, according to measures η and γ. See Section 7 for details.

Estimating the measures. The measures we consider are all based on the expected values of all possible deterministic memoryless strategies in the game. Since the state space of most of the games we consider in this section is extremely large, two levels of sampling were necessary: first, it is not possible to evaluate all strategies; second, it is not even possible to evaluate all possible paths of a single strategy. Both levels were adressed by uniform random sampling. As a rule of thumb, we chose the sizes of the samples in such a way that approximately one hour of computation is devoted to each game, on an Intel Core i5, 3.3Ghz and with a sequential implementation.

The game Tic-tac-toe is an exception, as it allows less than a thousand different strategies and 400 paths for each strategy. Hence, it was possible to analyze it exhaustively. For the other games, the sampling introduces estimation errors whose precise characterization goes beyond the scope of this paper. Indeed, our purpose is to explore and compare different notions of controllability, and the accuracy achieved by the experiments is sufficient to clearly distinguish how different measures evaluate different games.

We now briefly present the games that we have evaluated and discuss the corresponding experimental results. Since our framework is tailored to solitaire games with a Boolean outcome, most games have been modified to fit these constraints.

Tic-tac-toe. Tic-tac-toe is a well-known children game, played on a 3x3 board by two players. We derive 5 different solitaire versions of this game, by fixing the strategy of Player 2 in one of the following ways:

1. Put the token uniformly at random (in one of the available cells).
2. If there is a move that leads to immediate victory of Player 2, play that move, otherwise play uniformly at random.
3. If there is a move that prevents immediate victory of Player 1, play that move, otherwise play uniformly at random.
4. Avoid an immediate loss, or else try to obtain an immediate victory, or else play uniformly at random.
5. Put the token in the first available cell (deterministic strategy).

Moreover, a tie is considered a victory for Player 2. In Figure 2, we refer to Tic-tac-toe with strategy i (for $i = 1, \ldots, 5$) as "TTT-i".

Video Poker. Video poker is a very common solitaire casino game, similar to a slot machine, but based on the basic rules of poker. In our version, at each round the player pays a fixed bet of 1 dollar and is dealt 5 cards from a shuffled deck of 52 cards. Then, he can choose which cards to keep and which to change. Once the discarded cards have been redealt from the same deck, a payoff table associates payments to combinations of cards. If the player has a positive balance, another round is played with a newly shuffled deck of cards. To obtain a Boolean game, we fix two parameters: n is the initial balance in dollars, and k is the number of rounds. The player starts with an initial balance of n dollars and wins if he keeps a positive balance for k rounds. In Figure 2, we refer to the parameterized game as "VP-(n, k)".

Black Jack. Black Jack is a game usually played by several players against a dealer, with no direct interaction between players (except for card counting purposes). The rules of the game can easily be found online. In our simplified version, the player initially pays a fixed bet of 2 dollars and wins back the bet if he scores higher than the dealer (but less than 21) and 3 dollars if he wins with Black Jack (21). Then, if the player has a positive balance, another round is played using the same stack of cards. Hence, contrary to video poker, in the rounds after the first one the player can use a card counting strategy to improve his odds. The player wins if she keeps a positive balance for k rounds, where k is a parameter. In Figure 2, we refer to this game as "BJ-k".

Minesweeper. Our version of Minesweeper is played on a 5x5 grid, on which n mines are randomly positioned at the beginning of each play. The player, who begins with no information on the position of the mines, in each round selects a cell and is told how many mines are located in the 8 adjacent cells. The game continues until either the player chooses a cell containing a mine or all empty cells have been uncovered. In Figure 2, we refer to the game with n mines as "MS-n".

It is intuitively clear that it is very hard to win at Minesweeper by playing at random. Hence, random strategies have very small expected values, while comparatively few strategies have significant chances to win [2].

Archery. We include in the comparison an individual (i.e., solitaire) sport, as an example of human activity that is unanimously considered skill-driven. Archery fits well this role, also because its outcomes are precisely quantifiable with a fine grain, if one is able to measure the distance of each single shot from the center of the target. This is precisely what was done in a recent paper by Kolayiş et al. [15], whose purpose was to statistically characterize the typical spatial distribution of arrows. To this aim, 18 athletes were asked to throw up to 12 arrows each, and the position of each arrow from the target center was measured using digital videography. For evaluating this "game", we consider each athlete as a strategy and the distance of each arrow from the target center as an outcome of that strategy. To turn outcomes into Boolean values, we consider each shot a "win" if it is closer to the center of the target than the average shot. We can then apply our measures to the data because they depend only on the expected values of each strategy.

7.1 Discussion

Figure 2 displays a selection of our experimental results, on the same logarithmic scale. But for a few exceptions, both measures agree on the order in which different games should be put according to player controllability.

In the case of Tic-tac-toe, both measures confirm the intuition that the 5 variants exhibit a decreasing amount of randomness, with variant 1 being the most random. In particular, both measures assign value 1 to version 5 of the game, which is entirely deterministic. All Tic-tac-toe variants are significantly more controllable than the two card games we examine.

[2] Becerra [1] reports an optimal expected value of at least 0.91 for a board of 9x9 with 10 mines.

In particular, a short game of Black Jack achieves approximately 1% controllability according to both measures, while Video Poker scores even lower.

In analyzing Minesweeper, the measures exhibit a rather extreme case of non-co-monotonicity. Indeed, when increasing the number of mines, η approaches a positive value around $7 \cdot 10^{-3}$, whereas γ approaches 1. We conjecture that this contrast is tied to the extremely low average probability of winning in this game (with 6 mines, the probability of winning by playing at random is about $1.6 \cdot 10^{-5}$).

As expected, archery sits near the top of our ranking, achieving measures $\eta = 0.47$ and $\gamma = 0.38$. With respect to the "preponderance test" mentioned in Section 4, we observe that setting the threshold between games of skill and games of chance at $\frac{1}{2}$ would put even archery in the first class, albeit by a thin margin [3].

8. CONCLUSIONS

The balance between skill and luck in games is a topic of interest for players, game designers, lawmakers, and entrepreneurs in the gambling business. In this paper, we put forward a formal framework to evaluate such balance on 1-player games, general enough to embrace a wide variety of games, both existing and yet-to-come. As a consequence, it is the first time, to the best of our knowledge, that common games ranging from Tic-tac-toe to Black Jack have been compared automatically and on equal footing.

The theoretical and experimental results also show the limits of this preliminary investigation. Both Proposition 5 and the experiments suggest that the proposed axioms are compatible with widely differing weak orders. Hence, it may be desirable to strengthen them to approach a full axiomatic characterization of a single notion of controllability. For instance, all our axioms and measures are symmetrical w.r.t. winning and losing. Breaking this symmetry may lead to measures that are closer in spirit to the common understanding of what constitutes "control" in a game.

Finally, a natural development would be to extend the present investigation to multi-player games. However, that setting poses radically different challenges, starting from the fact that different players may have different control capabilities on the game, as the game needs not be symmetric w.r.t. the players.

Acknowledgments

Clara Pappalardo and Ugo Pugliese originally brought the topic of this paper to the attention of the author. Christos Stergiou took part in preliminary discussions and contributed to the experimental setup described in Section 7. The evaluation of archery was made possible by the kind cooperation of Ipek Kolayiş and her co-authors [15], who agreed to share the data concerning their experiment.

REFERENCES

[1] D. Becerra. Algorithmic approaches to playing Minesweeper, 2015. Bachelor's thesis, Harvard College.

[2] P. Borm and B. van der Genugten. On a relative measure of skill for games with chance elements.

[3] Our observations about archery apply to a specific set of data which does not necessarily represent a random sample of athletes.

Trabajos de Investigaciòn Operativa (Top), 9(1):91–114, 2001.

[3] P. Cohen. *Empirical Methods for Artificial Intelligence*. MIT Press, Cambridge, MA, USA, 1995.

[4] J. K. Dagsvik. Choice under Uncertainty and Bounded Rationality. Discussion Papers 409, Statistics Norway, Research Department, Mar. 2005.

[5] M. Dreef, P. Borm, and B. van der Genugten. Measuring skill in games: several approaches discussed. *Mathematical Methods of Operations Research*, 59(3):375–391, 2004.

[6] M. Dreef, P. Borm, and B. van der Genugten. A new relative skill measure for games with chance elements. *Managerial and Decision Economics*, 25(5):255–264, 2004.

[7] E. Feinberg and A. Shwartz, editors. *Handbook of Markov Decision Processes: Methods and Applications*. Springer, 2002.

[8] G. Gamst, L. Meyers, and A. Guarino. *Analysis of Variance Designs*. Cambridge University Press, 2008.

[9] L. Hedhes and I. Olkin. *Statistical Methods for Meta-Analysis*. Academic Press, 1985.

[10] S. Heubeck. Measuring skill in games: a critical review of methodologies. *Gaming Law Review and Economics*, 12(3):231–238, 2008.

[11] S. Heubeck. Measuring skill in games: Evaluating legality. *Review of Law and Economics*, 4(1):25–34, 2008.

[12] S. Johnson. Playing the odds. In C. Steinkuehler, K. Squire, and S. Barab, editors, *Games, Learning, and Society*. Cambridge Univ. Press, 2012.

[13] J. Kadane. *Statistics in the law*. Oxford Univ. Press, 2008.

[14] R. Kaye. Minesweeper is NP-complete. *The Mathematical Intelligencer*, 22(2):9–15, 2000.

[15] I. Kolayiş, M. Çilli, H. Ertan, and J. Knicker. Assessment of target performance in archery. *Procedia - Social and Behavioral Sciences*, 152:451 – 456, 2014. ERPA International Congress on Education, Istanbul, Turkey.

[16] S. Levitt and T. Miles. The role of skill versus luck in poker evidence from the world series of poker. *Journal of Sports Economics*, 15(1):31–44, 2014.

[17] J. McCrory. Video poker and the skill versus chance debate. *Gaming Law Review*, 6(3):223–227, 2002.

[18] S. Olejnik and J. Algina. Measures of effect size for comparative studies: Applications, interpretations, and limitations. *Contemporary Educational Psychology*, 25(3):241 – 286, 2000.

[19] M. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.

[20] R. Potter van Loon, M. van den Assem, and D. van Dolder. Beyond chance? The persistence of performance in online poker. *PLoS ONE*, 10(3):e0115479, 2015.

[21] B. Ratitch and D. Precup. Characterizing Markov decision processes. In *Proceedings of the 13th European Conference on Machine Learning*, ECML '02, pages 391–404, London, UK, UK, 2002. Springer-Verlag.

[22] A. Wyner. Can the skill level of a game of chance be measured? *Chance*, 25(3):50–53, 2012.