

Variance Decompositions for Extensive-Form Games

Alex Cloud
North Carolina State University
Raleigh, United States
kacloud@ncsu.edu

Eric B. Laber
Duke University
Durham, United States
eric.laber@duke.edu

Abstract—The extent to which an individual or chance can influence the outcome of a game is a central question in the analysis of games. Consequently, the ability to characterize sources of variation in game outcomes may have significant implications in areas such as game design, law, and multi-agent reinforcement learning. We derive a closed-form expression and estimators for the variance in the outcome of a general multi-agent game that is attributable to a player or chance. We analyze poker hands to show that randomness in the cards dealt has surprisingly little influence on the outcomes of each hand. A simple example is given that demonstrates how variance decompositions can be used to measure other interesting properties of games.

Index Terms—luck and skill in games, descriptive statistics, variance decomposition

I. INTRODUCTION

From game design studios to courtrooms, randomness in games has been the subject of extensive discussion. Game designers use random game elements to protect players’ egos, increase gameplay variety, and limit the efficacy of mental calculation [1]. In U.S. state law, the question of whether Poker is predominantly a game of chance or skill is considered to be central to the legality of online Poker [2], [3].

The question of how to measure the role of luck versus skill has proved difficult and produced many answers [1], [3]–[8]. For example, in *USA v. Lawrence Dicristina*, economic consultant and high-level amateur poker player Randal Heeb testified that “statistical analysis of poker hands confirms that skill predominates over chance.” His conclusion was based on a series of heuristic data analyses combined with intuitive judgments [9]. Others have argued that the strong association between player skill rating and future earnings constitute strong evidence that poker should be considered a game of skill [3], [10].

A first step in assessing the role of chance in a game is to quantify sources of uncertainty. We examine how variation in the outcomes of a game can be attributed to players or chance events by expressing variation in game outcomes as the sum of variance components associated with (i) the actions taken by a player of interest, and (ii) all remaining sources of variation. By applying this decomposition to a conceptual “chance player,” we measure the degree to which randomness inherent in a game biases the results in favor of a given player.

We derive an analytical expression for these variance components and use it to obtain estimators which are model-free in the sense that they do not require access to an entire game model or other players’ behavior. Our results apply to finite extensive-form games in general. As an illustrative example, we analyze poker hands played by the DeepStack poker agent against professional players [11] and find that chance events have very little influence on the expected per-hand profit for a player relative to the total variation in per-hand profit. A roadmap of the paper is as follows:

- Section II defines extensive-form games.
- Section III-A casts a general extensive-form game in terms of random variables. Section III-B describes how to decompose the variance of a game outcome into the sum of two nonnegative terms, gives a formula for the terms, and provides two ways of estimating them.
- Section IV gives estimates for the variance component for chance in collections of Poker games.
- Section V offers an idea for a further variance decomposition for measuring skill, chance, and non-transitivity in games and applies it to a conceptual game.
- Section VI discusses the interpretation and relevance of the results.

II. EXTENSIVE-FORM GAMES

An *extensive-form game* is a tree-based representation of a multi-agent system; Figure 1 displays a simple example. In this representation, the game is played by traversing the tree from the root to a leaf node, with a player’s action at each node determining the next node visited. Our notation is based on [12] and [13], with some modifications.

Let \mathcal{S} denote the set of possible game states which we assume is finite; each state is associated with a node in the game tree. Define $\mathcal{N} = \{1, \dots, n\}$ to be the set of (non-chance) players and let c denote the chance player. The *player function* $P : \mathcal{S} \rightarrow \mathcal{N} \cup \{c\}$ associates each state with a player. At each state $s \in \mathcal{S}$, there are a finite number of available actions $\mathcal{A}(s)$, such that each $a \in \mathcal{A}(s)$ uniquely determines the next state visited in the tree [14].

A sequence of actions $z = (a_1, \dots, a_m)$ is a *terminal history* if it leads from the root to a leaf of the game tree; let \mathcal{Z} denote the set of all terminal histories. For each player $i \in \mathcal{N}$ and terminal history $z \in \mathcal{Z}$, a reward $r^i(z) \in \mathbb{R}$ is obtained




Fig. 1. An example of an extensive-form game. Each node in the tree is a state $s \in \mathcal{S}$ and is annotated with the corresponding player, $P(s)$. The dashed line represents Player 2's information state; in this example, they cannot tell what move Player 1 played. Rewards for Player 1 are shown below the terminal nodes.

by player i upon reaching z . Each player $i \in \mathcal{N}$ has a set of *information states* \mathcal{U}^i which represent collections of nodes which are indistinguishable to the player. In particular, \mathcal{U}^i is a partition of $\{s \in \mathcal{S} : P(s) = i\}$ with the additional condition that $\mathcal{A}(s) = \mathcal{A}(s')$ if s and s' are in the same information state. So, we can write $\mathcal{A}(u)$ for $u \in \mathcal{U}^i$ unambiguously. Define $\mathcal{U}^c = \{\{s\} : P(s) = c\}$. We consider games of *perfect recall*, so that for every player i , each $u \in \mathcal{U}^i$ can be uniquely identified with the sequence of information states and actions required to arrive there.

Finally, the behavior of each player $i \in \mathcal{N} \cup \{c\}$ is described by a *policy* π^i (also known as a behavioral strategy), which is a function that maps each information state $u \in \mathcal{U}^i$ to a distribution over the allowable actions $\mathcal{A}(u)$. A *policy profile* is a tuple of player policies, $\pi = (\pi^1, \dots, \pi^n)$. By convention, the policy of the chance player π^c is considered to be a fixed part of the extensive-form game itself and not a part of any policy profile.

III. VARIANCE DECOMPOSITIONS FOR GAME OUTCOMES

To formalize our results, we represent an extensive-form game in terms of random variables. Having done this, the outcome of the game will be a random variable Y indicating the score obtained by a particular player.

A. Extensive-form games with random variables

First, we introduce random variables that represent the actions selected by players in a single play of the game. For each $i \in \mathcal{N} \cup \{c\}$, and for each $u \in \mathcal{U}^i$, let $A(u)$ be a random variable taking values in $\mathcal{A}(u)$ which represents the action player i would take given information state u . This variable always realizes a value, even if u is not reached in a particular play of the game. Note that for $u \neq u' \in \mathcal{U}^i$, it need not be the case that $A(u)$ is independent of $A(u')$. This manner of specifying player behavior is quite general and can account for different models of player action selection. For example, a player may randomly precommit to a deterministic policy (this is known as a mixed strategy in the game theory literature), or select actions independently at random at each time step (a behavioral strategy) [15].

For each terminal history $z \in \mathcal{Z}$ and player $i \in \mathcal{N} \cup \{c\}$, let $m^i(z)$ be the number of actions selected by player i along z , so that for each $j \in \{1, \dots, m^i(z)\}$, we can write $u_{z,j}^i$ and $a_{z,j}^i$ to denote the j th information state observed and action selected by player i along terminal history z . Define $I_{z,j}^i = \mathbb{1}[A(u_{z,j}^i) = a_{z,j}^i]$ to be the Bernoulli random variable that indicates whether player i selects $a_{z,j}^i$ at $u_{z,j}^i$. Finally, define $I_z^i = \prod_{j=1}^{m^i(z)} I_{z,j}^i$ to be the Bernoulli random variable that indicates whether player i selects all actions along z . (If $m^i(z) = 0$, set $I_z^i \equiv 1$.)

A terminal history occurs if and only if every action along it is selected. Therefore, for each $z \in \mathcal{Z}$, $I_z = \prod_{i \in \mathcal{N} \cup \{c\}} I_z^i$ defines a Bernoulli random variable such that the success probability $P(I_z = 1)$ is the probability that terminal history z is realized. Let Z be a random terminal history variable such that $P(Z = z) = P(I_z = 1)$ for all z that represents a random play-through of the game. This allows us to cast the outcome of an extensive-form game as

$$\mathbf{Y} = [r^1(Z), \dots, r^n(Z)].$$

Write $Y = r(Z) = r^h(Z)$, the random variable representing the reward for a player $h \in \mathcal{N}$ upon one play of the game. Our goal is to express its variance, $V(Y) = E\{[Y - E(Y)]^2\}$, as a sum of nonnegative terms corresponding to meaningful properties of the game.

B. Variance decomposition

Let $i \in \mathcal{N} \cup \{c\}$ be a player of interest, and let $\mathbf{A}^i = [A(u)]_{u \in \mathcal{U}^i}$ be the concatenation of all actions for player i . By the law of total variance we can decompose the variance in game outcomes as

$$V(Y) = V[E(Y|\mathbf{A}^i)] + E[V(Y|\mathbf{A}^i)]. \quad (1)$$

The term $E(Y|\mathbf{A}^i)$ is the average game outcome upon many traversals of the game tree when player i commits ahead of time to playing the actions in \mathbf{A}^i . For example, $E(Y|\mathbf{A}^c)$ represents the average outcome for a group of poker players who play the same hand from a deck with a particular card order many times, or the average outcome for a pair of chess players who start with the same colors every game. Then $V[E(Y|\mathbf{A}^c)]$ is the variation in this mean as the chance actions \mathbf{A}^c vary, and represents the variation in game outcomes “explained by” chance events. The latter term of (1) has a similar interpretation as the variation in game outcomes not explained by actions selected by player i .

Let $i \in \mathcal{N} \cup \{c\}$ be a player of interest. Suppose that player i plays according to a behavioral strategy π^i , meaning that $A(u)$ is independent of $A(u')$ for all $u \neq u' \in \mathcal{U}^i$ and action probabilities are given by a policy such that $P(A_{z,k}^i = 1) = \pi^i(a_{z,k}^i | u_{z,k}^i)$ for all $z \in \mathcal{Z}$ and $k \in \{1, \dots, m^i(z)\}$. No such assumption is required for the remaining players; we only require that other players' (precommitments to) actions are independent of the actions of player i .

Let $\eta^i(z) = P(I_z^i = 1)$ be the probability that player i assigns to actions along terminal history z ; similarly,

$\eta^{-i}(z) = P(\prod_{i' \in \mathcal{N} \cup \{c\} \setminus \{i\}} I_{z'}^{i'} = 1)$ is the probability assigned by other players, and $\eta(z) = \eta^i(z) \eta^{-i}(z) = P(I_z = 1)$. For each information state $u \in \mathcal{U}^i$, define $\mathcal{Z}(u) = \{z \in \mathcal{Z} : u \text{ is visited in } z\}$, $\eta(u) = \sum_{z \in \mathcal{Z}(u)} \eta(z)$, $\eta^i(u) = \sum_{z \in \mathcal{Z}(u)} \eta^i(z)$, and $\eta^{-i}(u)$ such that $\eta(u) = \eta^i(u) \eta^{-i}(u)$. For each $a \in \mathcal{A}(u)$, define $\mathcal{Z}(ua) = \{z \in \mathcal{Z} : u \text{ is visited in } z \text{ and action } a \text{ is selected at } u.\}$. Define $q(u, a) = E[r(Z)|Z \in \mathcal{Z}(ua)]$ to be the expected outcome given that player i is at u and takes action a ; similarly, define $v(u) = E[r(Z)|Z \in \mathcal{Z}(u)]$.

Our main result is an expression of the variance in game outcomes explained by player i 's actions as a sum of weighted, squared action-value and value functions over all of player i 's information states:

$$V[E(Y|\mathbf{A}^i)] = \sum_{u \in \mathcal{U}^i} \left(\sum_{a \in \mathcal{A}(u)} [q(u, a)]^2 \pi^i(a|u) - [v(u)]^2 \right) \eta^{-i}(u) \eta(u). \quad (2)$$

A proof is provided in Appendix A. Computing this requires traversing the game tree a fixed number of times and hence is $O(|\mathcal{S}|)$. From this we obtain a formula for the other variance component by observing that $E[V(Y|\mathbf{A}^i)] = V(Y) - V[E(Y|\mathbf{A}^i)]$, where $V(Y)$ can be evaluated as $\sum_{z \in \mathcal{Z}} [r(z) - \sum_{z' \in \mathcal{Z}} r(z') \eta(z')]^2 \eta(z)$.

Assuming η^{-i} , q , and v are known, given an i.i.d. sequence of ν playthroughs of the game, each generating a sequence $\bar{U}_k = (U_{k,1}, \dots, U_{k,l_k})$ of observed information states in \mathcal{U}^i , then the following is a consistent estimator for $V[E(Y|\mathbf{A}^i)]$ as proved in Appendix B:

$$\nu^{-1} \sum_{k=1}^{\nu} \sum_{l=1}^{l_k} \left(\sum_{a \in \mathcal{A}(U_{k,l})} [q(U_{k,l}, a)]^2 \pi^i(a|U_{k,l}) - [v(U_{k,l})]^2 \right) \eta^{-i}(U_{k,l}). \quad (3)$$

In practice, q and v can be estimated by supervised learning and $\eta^i = \eta/\eta^{-i}$ can be estimated with $\hat{\eta}(u) = \nu^{-1} \sum_{k=1}^{\nu} \sum_{l=1}^{l_k} \mathbb{1}(U_{k,l} = u)$ and $\eta^i(u) = \pi^i(u)$ (assuming the analyst does not have access to opponent policies and observations). However, if there are many possible information states, i.e., $|\mathcal{U}^i|$ is large, $\hat{\eta}(u)$ will greatly overestimate the visit probability. An alternative is a more straightforward regression-based estimator. Our regression-based estimator works by fitting a model for the conditional mean of the game outcome given a player's actions, then computing the empirical variance of the conditional mean estimator. The procedure is:

- 1) Specify a model $f_{\hat{\theta}}$ that maps the collection of all actions for the player of interest to a real number, $f_{\hat{\theta}} : \times_{u \in \mathcal{U}^i} \mathcal{A}(u) \rightarrow \mathbb{R}$.
- 2) For a each observed game $k \in \{1, \dots, \nu\}$, record action-outcome pairs (\mathbf{A}_k^i, Y_k) . For each k , if an information state for the player of interest, $u \in \mathcal{U}^i$ was not visited in game k , sample $A(u) \sim \pi^i(\cdot|u)$ and include the

sampled action in \mathbf{A}_k^i . Fit the model on the action-outcome pair data to find a $\hat{\theta}$ that minimizes the mean square error, $\nu^{-1} \sum_{k=1}^{\nu} [Y_k - f_{\hat{\theta}}(\mathbf{A}_k^i)]^2$, so $f_{\hat{\theta}}(\cdot)$ estimates $E(Y|\mathbf{A}^i = \cdot)$.

- 3) Estimate $V[E(Y|\mathbf{A}^i)]$ using the sample analog with $f_{\hat{\theta}}(\mathbf{A}^i)$ plugged in for $E(Y|\mathbf{A}^i)$:

$$\nu^{-1} \sum_{k=1}^{\nu} \left[f_{\hat{\theta}}(\mathbf{A}_k^i) - \nu^{-1} \sum_{h=1}^{\nu} f_{\hat{\theta}}(\mathbf{A}_h^i) \right]^2. \quad (4)$$

The regression-based estimator is consistent if $f_{\hat{\theta}}$ is consistent for $E(Y|\mathbf{A}^i = \cdot)$; a proof is provided in Appendix B.

IV. ANALYSIS OF PROFESSIONAL POKER PLAYERS VERSUS DEEPSTACK

We analyze 150 thousand hands of heads-up no-limit poker played by different players against the DeepStack poker agent, including 45 thousand hands played by self-identified professional players. For details on how the data were generated, see the supplemental materials of the DeepStack paper [11]. Our goal is to understand the role chance has in influencing the per-hand profits of a human playing against DeepStack, so we will estimate the variance component for chance for games played by each human player indexed by $j \in \{1, \dots, 33\}$. We also include an algorithm used for poker agent evaluation called Local Best Response (index $j = 0$), which we include as a form of transfer learning in order to improve estimates of expected outcomes for the human players. Assume that player j plays according to a policy π_j and write $E_{\pi_j}(Y|\mathbf{A}^c)$ to denote the expected per-hand profit for player j against DeepStack given all chance events \mathbf{A}^c . Then we would like to know $V[E_{\pi_j}(Y|\mathbf{A}^c)]$ for each j .

We use a neural network to estimate $E_{\pi_j}(Y|\mathbf{A}^c)$ given a player and the realization of all chance events:

- The player's pocket cards (2 cards)
- DeepStack's pocket cards (2 cards)
- The flop (3 cards)
- The turn (1 card)
- The river (1 card)

The neural network shares a representation of cards across all inputs: each card rank (e.g. Ace) and suit (e.g. hearts) is associated with a learned vector embedding; a card is represented by the concatenation of these embeddings. To capture the unordered nature of players' pocket cards and the flop, the card representations for each of those groups is summed. The architecture is depicted in Figure 2 and hyperparameters are given in Appendix C.

Our model was trained by stochastic gradient descent with the Adam optimizer [16] with early stopping based on cross-validation loss using a 90%/10% train-test split. If a hand ended before all chance events were observed (for example, if a player folded before the river), the cards associated with that chance event were randomly sampled from the remaining cards in the deck at that point in the game. These cards were resampled in each epoch of training in order to decrease variance. We present our results in Table I. For each player, the

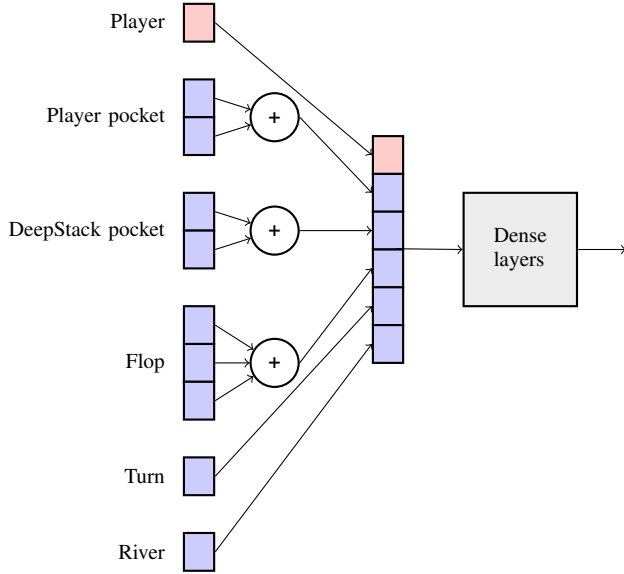


Fig. 2. The neural network architecture used for analysis of DeepStack hands. The input for each card (shown in blue) is a concatenation of the rank and suit of the card. The rank and suit are each assigned a vector embedding, with the same weights shared for all card inputs.

empirical variance of the regression estimator was computed over both the training and test data and is recorded in column “Chance var.” Due to randomness in the training procedure (neural network initialization, train-test split, and sampled actions), we repeat the procedure 100 times and report average results and standard deviations.

The results are somewhat surprising: typical values for the percent of total variance “explained” by chance events fall between 0% and 5%, with low standard deviations. We conclude that the influence of chance events alone on per-hand outcomes is quite limited. Rather, the large amount of variation in per-hand profits is mostly explained by player randomization and the interaction between those actions and chance. We elaborate in the Discussion section.

V. A THREE-WAY DECOMPOSITION FOR ASSESSING SKILLFULNESS OF A GAME

As another example of using variance decompositions to analyze games, we present a concept for measuring skill, chance, and non-transitivity that is inspired by prior work on decompositions of games [17] and recent developments regarding learning in the context of complex games with nontransitive elements [18] [19]. For simplicity, assume we are given a symmetric two-player zero-sum game and a population of players represented by a finite set of policies Π , each with a skill rating ρ_π for $\pi \in \Pi$. One notion of the skillfulness of the game is the variance in outcomes explained by players’ skill ratings alone, assuming two policies (π_1, π_2) are sampled uniformly from Π :

$$V(Y) = V[E(Y|\rho_{\pi_1}, \rho_{\pi_2})] + E[V(Y|\rho_{\pi_1}, \rho_{\pi_2})]. \quad (5)$$

Applying the law of total variance to the conditional variance $V(Y|\rho_{\pi_1}, \rho_{\pi_2})$, we condition on chance actions \mathbf{A}^c as in (1)

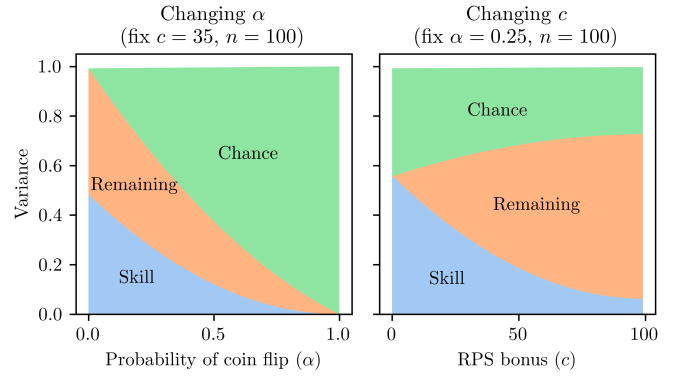


Fig. 3. Three-way variance decompositions for SkillRPS with different game parameters under the assumption that players selects moves independently and uniformly at random, i.e., for $i \in \{1, 2\}$, $N_i \sim \text{Uniform}(\{1, \dots, n\})$ and $A_i \sim \text{Uniform}(\{\text{Rock, Paper, Scissors}\})$ and are independent. Details on the variance components for SkillRPS are included in Appendix D.

to obtain $V(Y|\rho_{\pi_1}, \rho_{\pi_2}) = V[E(Y|\mathbf{A}^c, \rho_{\pi_1}, \rho_{\pi_2})|\rho_{\pi_1}, \rho_{\pi_2}] + E[V(Y|\mathbf{A}^c, \rho_{\pi_1}, \rho_{\pi_2})|\rho_{\pi_1}, \rho_{\pi_2}]$. Using linearity of expectation and the tower rule, this allows us to extend (5) to

$$\begin{aligned} V(Y) &= V[E(Y|\rho_{\pi_1}, \rho_{\pi_2})] && \text{(skill)} \\ &+ E\{V[E(Y|\mathbf{A}^c, \rho_{\pi_1}, \rho_{\pi_2})|\rho_{\pi_1}, \rho_{\pi_2}]\} && \text{(chance)} \\ &+ E[V(Y|\mathbf{A}^c, \rho_{\pi_1}, \rho_{\pi_2})]. && \text{(remaining)} \end{aligned} \quad (6)$$

We apply this formula to analyze a simple game parametrized by constants $n \in \mathbb{N}$, $c \in \mathbb{N} \cup \{0\}$, and $\alpha \in [0, 1]$ that is an abstract model of a game with a skill component (some strategies are strictly better than others), a nontransitive component (there exist cycles of pure strategies), and chance (some games are decided by events entirely out of the players’ hands). Skillful Rock Paper Scissors, or SkillRPS(n, c, α), is defined as follows: each player $i \in \{1, 2\}$ simultaneously selects a number $N_i \in \{1, \dots, n\}$ and a move $A_i \in \{\text{Rock, Paper, Scissors}\}$. Player 1’s score is $S = N_1 - N_2 + c \cdot \text{RPS}(A_1, A_2)$, where RPS is the payoff function for Rock Paper Scissors depicted in Table II.

The outcome of the game for player 1 is $Y = (1 - W)[\mathbb{1}(S > 0) - \mathbb{1}(S < 0)] + W(2Z - 1)$, where $W \sim \text{Bernoulli}(\alpha)$ and $Z \sim \text{Bernoulli}(1/2)$ are chance events such that W determines whether the game is decided by a fair coin flip Z . Note that when $n = 1$, $c > 0$, $\alpha = 0$, the game is classic Rock Paper Scissors, when $\alpha = 1$ it is a coin flip, and when $c = 0$ it is a transitive game. The game can be represented in extensive form as shown in Figure 1, which depicts SkillRPS(2, 0, 0.5).

In Figure 3, the three-way decomposition is given across many values of the SkillRPS game parameters, showing that the components correspond to meaningful properties of games: increasing the probability that the game outcome is determined by a coin flip increases the chance variance component to 1 as the other variance components decrease smoothly; increasing the bonus for winning at Rock Paper Scissors decreases the skill component. In this case, the “remaining” variation

Player name	Hands played	Mean profit	Variance	Chance var.	Chance var. %
Local best response	106,221	-0.07	4.71	0.11 (0.04)	2.4 (0.9)
Ivan Shabalin	3,122	-0.03	3.42	0.09 (0.03)	2.7 (1.0)
Pol Dmit	3,026	-0.09	4.99	0.11 (0.05)	2.2 (0.9)
Muskan Sethi	3,010	-0.21	8.07	0.12 (0.05)	1.5 (0.7)
Dmitry Lesnoy	3,007	0.01	4.42	0.09 (0.03)	1.9 (0.7)
Stanislav Voloshin	3,006	0.01	3.27	0.11 (0.05)	3.3 (1.4)
Lucas Schaumann	3,004	-0.02	2.59	0.11 (0.04)	4.1 (1.5)
Phil Laak	3,003	-0.08	3.58	0.10 (0.04)	2.8 (1.1)
Antonio Parlavecchio	3,003	-0.11	7.22	0.13 (0.05)	1.8 (0.7)
Kaishi Sun	3,002	-0.00	4.14	0.11 (0.04)	2.6 (1.0)
Martin Sturc	3,001	0.05	2.58	0.09 (0.03)	3.5 (1.3)
Prakshat Shrimankar	3,001	-0.02	3.47	0.10 (0.04)	2.9 (1.2)
Tsuneaki Takeda	1,901	0.03	7.46	0.09 (0.04)	1.3 (0.6)
Youwei Qin	1,759	-0.20	14.80	0.11 (0.04)	0.7 (0.3)
Fintan Gavin	1,555	0.00	10.97	0.11 (0.05)	1.0 (0.4)
Giedrius Talacka	1,514	-0.05	11.46	0.12 (0.05)	1.0 (0.5)
Juergen Bachmann	1,088	-0.18	7.80	0.17 (0.09)	2.2 (1.2)
Sergey Indenok	852	-0.03	13.90	0.11 (0.05)	0.8 (0.3)
Sebastian Schwab	516	-0.18	6.25	0.10 (0.05)	1.7 (0.8)
Dara O'kearney	456	-0.02	3.37	0.15 (0.06)	4.5 (1.9)
Roman Shaposhnikov	330	0.09	3.95	0.09 (0.04)	2.3 (1.0)
Shai Zurr	330	-0.12	4.15	0.09 (0.04)	2.2 (0.8)
Luca Moschitta	328	-0.14	4.83	0.11 (0.07)	2.4 (1.4)
Stas Tishekvich	295	0.03	3.90	0.11 (0.05)	2.9 (1.2)
Eyal Eshkar	191	-0.07	8.77	0.13 (0.05)	1.5 (0.6)
Jefri Islam	176	-0.38	10.56	0.10 (0.05)	0.9 (0.4)
Fan Sun	122	0.13	9.27	0.12 (0.08)	1.2 (0.9)
Igor Naumenko	102	-0.09	0.61	0.08 (0.04)	13.1 (6.4)
Silvio Pizzarello	90	-0.51	10.44	0.17 (0.19)	1.7 (1.9)
Gaia Freire	76	-0.01	0.09	0.10 (0.06)	111.8 (63.9)
Alexander Bös	74	-0.00	1.29	0.05 (0.03)	3.9 (2.1)
Victor Santos	58	0.18	0.96	0.12 (0.07)	12.6 (7.5)
Mike Phan	32	1.12	25.58	0.07 (0.05)	0.3 (0.2)
Juan-Manuel Pastor	7	-0.73	1.14	0.06 (0.08)	5.7 (7.3)

TABLE I

ANALYSIS OF THE VARIANCE IN PER-HAND PLAYER PROFIT (IN \$1,000'S) FOR HUMAN PROFESSIONALS AGAINST THE DEEPSTACK POKER AGENT. STANDARD DEVIATIONS ARE GIVEN IN PARENTHESES AND ARE BASED ON 100 REPLICATIONS OF THE TRAINING PROCEDURE.

$a_1 \backslash a_2$	Rock	Paper	Scissors
Rock	0	-1	1
Paper	1	0	-1
Scissors	-1	1	0

TABLE II

THE PAYOFF FUNCTION $RPS(a_1, a_2)$.

corresponds directly to the non-transitivity introduced by the RPS component of the game.

VI. DISCUSSION

One might hope that the variance component for chance $V[E(Y|\mathbf{A}^c)]$ measures how lucky a game is in the context of the players playing the game. We argue that this is not the case, and conclude with thoughts on the applicability of variance component estimation for the analysis of games.

First, the variance component for chance does not measure how lucky a game is because by design it avoids measuring variation introduced by random player actions. Consider the classic version of Rock Paper Scissors (RPS) depicted in Table II. A cautious player can guarantee an expected payoff of 0 by assigning uniform probability to each action, causing the outcome of the game to be uniformly random over $\{-1, 0, 1\}$. For this reason, it is natural to view RPS as a game of luck— however, RPS as typically modeled does not have a chance

player. All variation in RPS comes from randomness in player action selection. So, if we are to call RPS a game of luck, then a notion of luck that only considers chance events is inadequate.

Second, the variance component for chance is conservative in that it only measures the marginal (average) effect of chance actions on game outcomes. It does not capture the interaction between chance events and player actions. For example, consider a variant of RPS in which one of the players is replaced with a chance player. If the non-chance player employs a uniform random policy, then the expected outcome is 0 regardless of action is selected by chance. Thus $E(Y|\mathbf{A}^c = a) = 0$ for each $a \in \{\text{Rock, Paper, Scissors}\}$. This means that for any chance policy, the variance component for chance is 0, yet from the player's perspective, against a uniform chance policy, it is as though the game outcome is entirely determined by chance!

What the variance component for chance actually measures is the per-game amount that chance biases the outcome in favor of a player. In both the examples given above, luck plays a significant role in the game outcomes, but the realization of chance events alone does not tend to significantly tilt the game in the favor of either player— so our measure evaluates to 0. Returning to the analysis of DeepStack poker hands, we see that despite the large amount of variation in per-hand profits

(of which any one realization could be called “lucky”) the game (as played at a high level) is in some sense fair: on a hand-by-hand basis, the average amount that the random deck order advantages or disadvantages a particular player is small.

Video game designers may find the variance component for chance helpful in assessing the per-play advantage gleaned by a player due to chance events. We speculate that for a rewarding game experience, the variance component should be kept low, or else players will feel a sense of limited agency. Returning to the question of the legality of poker, our measure could represent a sufficient (but not necessary) criterion for determining that a game is “predominantly due to chance:” if the ratio of the variance component for the chance player to the total variation is greater than 50%, then clearly the game outcomes could be said to be predominantly due to chance. The three-way variance decomposition in (6) offers a way to characterize meaningful properties of games that arise in the context of multiagent reinforcement learning and presents new research challenges such as (i) accounting for estimation error in the skill rating (however it is defined), and (ii) accounting for the actual distribution from which policies are sampled to play each other, which is often not uniform but rather skill-based, such that players with nearby skill ratings are likely to be placed together.

ACKNOWLEDGEMENT

We are grateful to members of Laber Labs, the AI Accelerator at Riot Games, and Mara Bandt-Law for input on early drafts and helpful discussions. We thank three anonymous reviewers for engaging thoughtfully with the paper and providing constructive feedback.

REFERENCES

- [1] G. S. Elias, R. Garfield, and K. R. Gutschera, *Characteristics of games*. MIT Press, 2012.
- [2] J. M. Kelly, Z. Dhar, and T. Verbiest, “Poker and the law: is it a game of skill or chance and legally does it matter?” *Gaming law review*, vol. 11, no. 3, pp. 190–202, 2007.
- [3] S. D. Levitt and T. J. Miles, “The role of skill versus luck in poker evidence from the world series of poker,” *Journal of Sports Economics*, vol. 15, no. 1, pp. 31–44, 2014.
- [4] M. A. DeDonno and D. K. Detterman, “Poker is a skill,” *Gaming Law Review*, vol. 12, no. 1, pp. 31–36, 2008.
- [5] R. Croson, P. Fishman, and D. G. Pope, “Poker superstars: Skill or luck? similarities between golf—thought to be a game of skill—and poker,” *Chance*, vol. 21, no. 4, pp. 25–28, 2008.
- [6] S. Heubeck, “Measuring skill in games: A critical review of methodologies,” *Gaming Law Review and Economics*, vol. 12, no. 3, pp. 231–238, 2008.
- [7] —, “Measuring skill in games with random payoffs: Evaluating legality,” *Review of Law & Economics*, vol. 4, no. 1, pp. 25–34, 2008.
- [8] D. Getty, H. Li, M. Yano, C. Gao, and A. Hosoi, “Luck and the law: Quantifying chance in fantasy sports and other contests,” *SIAM Review*, vol. 60, no. 4, pp. 869–887, 2018.
- [9] R. D. Heeb, “Report of randall d. heeb, phd (united states of american against lawrence discristina),” Case1:11-cr-00414 Document 77-1, 7 2012.
- [10] R. J. P. van Loon, M. J. van den Assem, and D. van Dolder, “Beyond chance? the persistence of performance in online poker,” *PLoS one*, vol. 10, no. 3, 2015.
- [11] M. Moravcik, M. Schmid, N. Burch, V. Lisý, D. Morrill, N. Bard, T. Davis, K. Waugh, M. Johanson, and M. Bowling, “Deepstack: Expert-level artificial intelligence in heads-up no-limit poker,” *Science*, vol. 356, no. 6337, pp. 508–513, 2017.

- [12] M. Lanctot, K. Waugh, M. Zinkevich, and M. Bowling, “Monte carlo sampling for regret minimization in extensive games,” in *Advances in neural information processing systems*, 2009, pp. 1078–1086.
- [13] J. Heinrich, M. Lanctot, and D. Silver, “Fictitious self-play in extensive-form games,” in *International Conference on Machine Learning*, 2015, pp. 805–813.
- [14] Y. Shoham and K. Leyton-Brown, *Multiagent systems: Algorithmic, game-theoretic, and logical foundations*. Cambridge University Press, 2008.
- [15] D. Koller and N. Megiddo, “Finding mixed strategies with small supports in extensive form games,” *International Journal of Game Theory*, vol. 25, no. 1, pp. 73–92, 1996.
- [16] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [17] O. Candogan, I. Menache, A. Ozdaglar, and P. A. Parrilo, “Flows and decompositions of games: Harmonic and potential games,” *Mathematics of Operations Research*, vol. 36, no. 3, pp. 474–503, 2011.
- [18] D. Balduzzi, M. Garnelo, Y. Bachrach, W. Czarnecki, J. Perolat, M. Jaderberg, and T. Graepel, “Open-ended learning in symmetric zero-sum games,” in *International Conference on Machine Learning*. PMLR, 2019, pp. 434–443.
- [19] S. Omidshafiei, K. Tuyls, W. M. Czarnecki, F. C. Santos, M. Rowland, J. Connor, D. Hennes, P. Muller, J. Pérolat, B. De Vylder *et al.*, “Navigating the landscape of multiplayer games,” *Nature communications*, vol. 11, no. 1, pp. 1–17, 2020.
- [20] A. W. Van der Vaart, *Asymptotic statistics*. Cambridge university press, 2000, vol. 3.

APPENDIX A

VARIANCE COMPONENT FORMULA DERIVATION

Here we derive (2), the formula for $V[E(Y|\mathbf{A}^i)]$. The basic strategy is to write the game outcome Y as a sum of random variables representing different paths through the game tree, then to use basic properties of probability and algebra to manipulate the expression, concluding with an inductive argument.

Recall that $I_z = \prod_{i \in \mathcal{N} \cup \{c\}} I_z^i = \prod_{i \in \mathcal{N} \cup \{c\}} \prod_{j=1}^{m^i(z)} I_{z,j}^i$ is the indicator that all actions along terminal history z are selected, $Y = \sum_{z \in \mathcal{Z}} r(z) I_z$, and $u_{z,j}^i$ is the j th information state observed by player i in terminal history z . Write $I_{z,k}^i = \prod_{j=k}^{m^i(z)} I_{z,j}^i$, the indicator that player i selects all actions in z at and after $u_{z,k}^i$. Let $d(u)$ be the depth of u in its trajectory; for example, if u is the first observation of a player in their trajectory, $d(u) = 1$. Define $W_u = \sum_{z \in \mathcal{Z}(u)} r(z) \eta^{-i}(z) I_{z,d(u)}^i$ and $\mathcal{U}_d^i = \{u \in \mathcal{U}^i : d(u) = d\}$. Then

$$\begin{aligned} V[E(Y|\mathbf{A}^i)] &= V\left[\sum_{z \in \mathcal{Z}} r(z) \eta^{-i}(z) I_z^i\right] \\ &= V\left[\sum_{u \in \mathcal{U}_1^i} \sum_{z \in \mathcal{Z}(u)} r(z) \eta^{-i}(z) I_z^i\right] \\ &= V\left(\sum_{u \in \mathcal{U}_1^i} W_u\right). \end{aligned} \quad (7)$$

Note that histories $z \in \mathcal{Z}$ that contain no information states for player i have $I_z^i \equiv 1$, so they are constant inside the conditional expectation, which is why the second and third expressions are equal.

By the perfect recall assumption, each information state u can be uniquely identified with the sequence of information states and actions required to reach u . Furthermore, the behavioral strategy assumption gives that $A(u)$ is independent

of $A(u')$ if $u \neq u' \in \mathcal{U}^i$. Therefore if $u_{z,j}^i \neq u_{z',j}^i$ for some j , then $I_{z,h}^i$ is independent of $I_{z',h'}^i$ for all $h, h' \in \{j, \dots, \min[m^i(z), m^i(z')]\}$. We conclude that W_u is independent of $W_{u'}$ if $u \neq u'$. This allows us to split up (7) as follows:

$$\begin{aligned} V[E(Y|\mathbf{A}^i)] &= V\left(\sum_{u \in \mathcal{U}_1^i} W_u\right) = \sum_{u \in \mathcal{U}_1^i} V(W_u) \\ &= \sum_{u \in \mathcal{U}_1^i} \left(V\{E[W_u|A(u)]\} + E\{V[W_u|A(u)]\} \right). \end{aligned} \quad (8)$$

The last equality holds by the law of total variance. To evaluate the components of (8), write $W_{ua} = \sum_{z \in \mathcal{Z}(ua)} r(z) \eta^{-i}(z) I_{z,[d(u)+1]}^i$ for each $a \in \mathcal{A}(u)$ so we have that

$$\begin{aligned} W_u &= \sum_{z \in \mathcal{Z}(u)} r(z) \eta^{-i}(z) I_{z,d(u)}^i \\ &= \sum_{a \in \mathcal{A}(u)} \sum_{z \in \mathcal{Z}(ua)} r(z) \eta^{-i}(z) I_{z,d(u)}^i I_{z,[d(u)+1]}^i \\ &= \sum_{a \in \mathcal{A}(u)} W_{ua} \mathbf{1}(A(u) = a). \end{aligned}$$

With this we obtain the following expressions for the components of (8):

$$\begin{aligned} V\{E[W_u|A(u)]\} &= \sum_{a \in \mathcal{A}(u)} E^2(W_{ua}) \pi^i(a|u) - \left[\sum_{a \in \mathcal{A}(u)} E(W_{ua}) \pi^i(a|u) \right]^2; \\ E\{V[W_u|A(u)]\} &= \sum_{a \in \mathcal{A}(u)} V[W_u|A(u) = a] P[A(u) = a]. \end{aligned}$$

Write $r(u, a) = E\{r(Z) \mathbf{1}[Z \in \mathcal{Z}(u, a)]\}$ and $r(u) = E\{r(Z) \mathbf{1}[Z \in \mathcal{Z}(u)]\}$. It can be shown that

$$E(W_{ua}) = r(u, a) [\eta^i(u) \pi(a|u)]^{-1}.$$

Substituting these terms back into the expression for the variance component, we get that

$$\begin{aligned} V\{E[W_u|A(u)]\} &= [\eta^i(u)]^{-2} \left(\sum_{a \in \mathcal{A}(u)} [r(u, a)]^2 \pi^i(a; u) - [r(u)]^2 \right); \\ E\{V[W_u|A(u)]\} &= \sum_{a \in \mathcal{A}(u)} V\left(\sum_{z \in \mathcal{Z}(ua)} r(z) \eta^{-i}(z) I_{z,[d(u)+1]}^i \right) \pi^i(a|u). \end{aligned}$$

Now take $Y_2 = \sum_{z \in \mathcal{Z}(ua)} r(z) \eta^{-i}(z) I_{z,[d(u)+1]}^i$ and repeat the steps shown in (8) inductively to obtain that:

$$\begin{aligned} V[E(Y|\mathbf{A}^i)] &= \sum_{u \in \mathcal{U}^i} \left(\sum_{a \in \mathcal{A}(u)} [r(u, a)]^2 / \pi^i(a|u) - [r(u)]^2 \right) / \eta^i(u). \end{aligned}$$

Because $r(u, a) = q(u, a) \eta(u) \pi^i(a|u)$ and $r(u) = v(u) \eta(u)$, this yields (2).

APPENDIX B CONSISTENCY PROOFS

First, we prove consistency of (3). Let $\mu(u) = \eta(u) / \sum_{u \in \mathcal{U}^i} \eta(u)$ be a normalized reach probability. Then

$$\begin{aligned} V[E(Y|\mathbf{A}^i)] &= \sum_{u \in \mathcal{U}^i} \left(\sum_{a \in \mathcal{A}(u)} [q(u, a)]^2 \pi^i(a|u) - [v(u)]^2 \right) \eta^{-i}(u) \eta(u) \\ &= \left(\sum_{u \in \mathcal{U}^i} \eta(u) \right) E_{U \sim \mu} \left[\left(\sum_{a \in \mathcal{A}(U)} [q(U, a)]^2 \pi^i(a|U) - [v(U)]^2 \right) \eta^{-i}(U) \right]. \end{aligned}$$

Note that

$$\begin{aligned} \sum_{u \in \mathcal{U}^i} \eta(u) &= \sum_{u \in \mathcal{U}^i} \sum_{z \in \mathcal{Z}(u)} \eta(z) = \sum_{u \in \mathcal{U}^i} \sum_{z \in \mathcal{Z}} \eta(z) \mathbf{1}(u \in z) \\ &= \sum_{z \in \mathcal{Z}} \eta(z) \sum_{u \in \mathcal{U}^i} \mathbf{1}(u \in z) = E[d^i(Z)], \end{aligned}$$

where $d^i(Z)$ is the length of the trajectory for player i in terminal history Z . So, by the law of large numbers, $\nu^{-1} \sum_{k=1}^{\nu} d^i(Z_k) \xrightarrow{\text{a.s.}} \sum_{u \in \mathcal{U}^i} \eta(u)$ as $\nu \rightarrow \infty$. Consider the Markov Chain $\{U_t\}_{t \in \mathbb{N}}$ defined by the information states for player i observed upon repeated independent playthroughs of the game and let $\phi(u) = \left\{ \sum_{a \in \mathcal{A}(u)} [q(u, a)]^2 \pi^i(a|u) - [v(u)]^2 \right\} \eta^{-i}(u)$. Then $T^{-1} \sum_{t=1}^T \phi(U_t) \xrightarrow{\text{a.s.}} E_{U \sim \mu}[\phi(U)]$ as $T \rightarrow \infty$ by a Law of Large Numbers for Markov Chains since $\{U_t\}$ is irreducible and positive recurrent.

Converting both these results to the notation of the original statement of the estimator, we have $\nu^{-1} \sum_{k=1}^{\nu} l_k \xrightarrow{\text{a.s.}} \sum_{u \in \mathcal{U}^i} \eta(u)$ and $(\sum_{k=1}^{\nu} l_k)^{-1} \sum_{k=1}^{\nu} \sum_{l=1}^{l_k} \phi(U_{k,l}) \xrightarrow{\text{a.s.}} E_{U \sim \mu}[\phi(U)]$ as $\nu \rightarrow \infty$. Therefore their product converges to the estimand, as desired.

To prove consistency of the regression-based estimator (4), we employ a basic style of argument from empirical process theory. For background on notation and concepts, see Chapter 19 of [20]. Let operator P denote expectation and operator \mathbb{P}_n denote empirical expectation. Assume without loss of generality that $E(Y) = 0$, so we can write $V[E(Y|\mathbf{A}^i)] = Pg$, where $g: \times_{u \in \mathcal{U}^i} \mathcal{A}(u) \rightarrow [-b, b]$ is given by $g(\mathbf{a}) = [E(Y|\mathbf{A}^i = \mathbf{a})]^2$, where the bound $b = \max_{z \in \mathcal{Z}} [r^i(z)]^2$ exists by finiteness of \mathcal{Z} . Let $\hat{g}_\nu = \min(f_{\hat{\theta}}^2, b)$, noting the implicit dependence of $\hat{\theta}$ on the sample size ν . Then $\mathbb{P}_n \hat{g}_\nu$ is our estimator and we have that

$$\begin{aligned} |\mathbb{P}_n \hat{g}_\nu - Pg| &= |\mathbb{P}_n \hat{g}_\nu - P\hat{g}_\nu + P\hat{g}_\nu - Pg| \\ &\leq |\mathbb{P}_n \hat{g}_\nu - P\hat{g}_\nu| + |P\hat{g}_\nu - Pg| \\ &\leq \sup_{g' \in \mathcal{G}} |\mathbb{P}_n g' - Pg'| + P|\hat{g}_\nu - g|, \end{aligned}$$

where \mathcal{G} refers to the set of bounded functions $g': \times_{u \in \mathcal{U}^i} \mathcal{A}(u) \rightarrow [-b, b]$. It is not hard to show that the finite domain and uniform boundedness of \mathcal{G} implies that its bracketing numbers are always finite. Thus, by a Glivenko-Cantelli theorem [20], the left-hand side converges to 0 in

probability. The right-hand side converges to 0 in probability if $\sup_{\mathbf{a}} |\hat{g}_v(\mathbf{a}) - g(\mathbf{a})| \xrightarrow{P} 0$, which by the continuous mapping theorem is equivalent to $f_{\hat{\theta}}$ being consistent (in a uniform sense) for $E(Y|\mathbf{A}^i = \cdot)$.

APPENDIX C NEURAL NETWORK HYPERPARAMETERS

- Embedding sizes: player - 3, card suit - 5, card rank - 14
- Card representation size: 10
- Dense layer sizes: 50, 20
- Activation functions: ReLU (hidden), identity (output)
- Batch size: 16
- Learning rate: 0.001
- Early stopping patience: 3 epochs

APPENDIX D SKILLRPS DECOMPOSITION DETAILS

In this section, we show how to derive the exact variance components for SkillRPS. Recall that in SkillRPS, the outcome is $Y = (1 - W)[\mathbf{1}(S > 0) - \mathbf{1}(S < 0)] + W(2Z - 1)$, where $S = N_1 - N_2 + c \cdot \text{RPS}(A_1, A_2)$. In this case, a player's selection of N_i is considered to indicate their skill level, and $\mathbf{A}^c = (W, Z)$ is the collection of all chance actions. Adapting the three-way decomposition equation (6) to SkillRPS yields

$$\begin{aligned} V(Y) &= V[E(Y|N_1, N_2)] && \text{(skill)} \\ &+ E\{V[E(Y|W, Z, N_1, N_2)|N_1, N_2]\} && \text{(chance)} \\ &+ E[V(Y|W, Z, N_1, N_2)]. && \text{(remaining)} \end{aligned}$$

Assuming that $N_1, N_2 \stackrel{\text{iid}}{\sim} \text{Uniform}(\{1, \dots, n\})$ and are independent of $A_1, A_2 \stackrel{\text{iid}}{\sim} \text{Uniform}(\{\text{Rock}, \text{Paper}, \text{Scissors}\})$, we can derive closed form expressions for each variance component.

Using routine probability manipulations, one can derive the following term for the variance in Y explained by the "skill" of the players in the case that the coin flip didn't happen ($W = 0$), for all $n \in \mathbb{N}$ and $c \in \mathbb{N} \cup \{0\}$. Begin by finding $E(Y|N_1 = n_1, N_2 = n_2, W = 0)$ for arbitrary $n_1, n_2 \in \{1, \dots, n\}$, which is easy since the only remaining source of variation is $\text{RPS}(A_1, A_2) \sim \text{Uniform}(\{-1, 0, 1\})$. Next, treat this term as a discrete random variable depending on N_1 and N_2 and compute its variance. This yields:

$$\begin{aligned} &V[E(Y|N_1, N_2, W = 0)] \\ &= \begin{cases} 1 - \frac{1}{n} & \text{if } c = 0 \\ 1 - \frac{1}{3n} + \frac{8c^2 + 2c - 16cn}{9n^2} & \text{if } 0 < c < n \\ (1 - \frac{1}{n})/9 & \text{if } c \geq n. \end{cases} \end{aligned}$$

Call this term $\psi(n, c)$. From here it can be shown that:

$$\begin{aligned} V[E(Y|N_1, N_2)] &= (1 - \alpha)^2 \psi(n, c) \\ E\{V[E(Y|W, Z, N_1, N_2)|N_1, N_2]\} &= \alpha + \alpha(1 - \alpha) \psi(n, c). \end{aligned}$$

Finally, it follows that

$$\begin{aligned} &E[V(Y|W, Z, N_1, N_2)] \\ &= \begin{cases} 0 & \text{if } c = 0 \\ (1 - \alpha) [1 - \frac{1}{n} + \frac{2c}{3n^2} - \psi(n, c)] & \text{if } 0 < c < n \\ (1 - \alpha) (\frac{8}{9} - \frac{2}{9n}) & \text{if } c \geq n. \end{cases} \end{aligned}$$