

Opponent Indifference in Rating Systems: A Theoretical Case for Sonas

Greg Bodwin ✉

University of Michigan, Ann Arbor, MI, United States

Forest Zhang¹ ✉

University of Michigan, Ann Arbor, MI, United States

Abstract

In competitive games, it is common to assign each player a real number *rating* signifying their skill level. A *rating system* is a procedure by which player ratings are adjusted upwards each time they win, or downwards each time they lose.

Many matchmaking systems give players some control over their opponent's rating; for example, a player might be able to selectively initiate games against opponents whose ratings are publicly visible, or abort a game without penalty before it begins but after glimpsing their opponent's rating. It is natural to ask whether one can design a rating system that does not incentivize a rating-maximizing player to act strategically, seeking games against opponents of one rating over another. We show the following:

- The *full* version of this “opponent indifference” property is unfortunately too strong to be feasible. Although it is satisfied by some rating systems, these systems lack certain desirable expressiveness properties, suggesting that they are not suitable to capture most games of interest.
- However, there is a natural relaxation, roughly requiring indifference between any two opponents who are both “reasonably evenly matched” with the choosing player. We prove that this relaxed variant of opponent indifference, which we call *P* opponent indifference, is viable. In fact, a certain strong version of *P* opponent indifference precisely characterizes the rating system *Sonas*, which was originally proposed for its empirical predictive accuracy on the outcomes of high-level chess games.

2012 ACM Subject Classification Theory of computation → Algorithmic game theory and mechanism design

Keywords and phrases Rating systems, opponent indifference, incentive compatibility, mechanism design, game theory

Digital Object Identifier 10.4230/LIPIcs.ITCS.2023.21

Acknowledgements We are grateful to Jeff Sonas for a helpful discussion on the history and relationship between rating systems, and we are grateful to an anonymous reviewer for thorough and helpful feedback.

1 Introduction

For two-player competitive games like chess, professional sports, online gaming, etc, it is common practice to assign players/teams a real-number *rating* capturing their skill level. Ratings are commonly applied in matchmaking systems [2, 3]; some more creative recent applications include pricing algorithms [26], item response theory [11], etc. (see Section 1.4 for more detail on the contribution of these papers). A *rating system* is an algorithm that adjusts a player's rating upwards after each win, or downwards after each loss. Some notable rating systems used in practice include Harkness [13], Elo [9], Glicko [12], Sonas [20], TrueSkill [14],

¹ Corresponding author



URS [1], and more. We refer to the papers [12, 18, 14], which each develop a notable rating system, for high-quality technical overviews of the inner workings and contrasts between the rating systems in the literature.

1.1 Our Model of Rating Systems

The exact model of rating systems studied in this paper is original, although really all of its ingredients are slight variants, rephrasings, or abstractions of ingredients that appear in prior work (we will address these similarities as we go). Our goal is to consider a maximally general model of rating systems that are:

- *One-dimensional*, meaning that they only track a rating parameter for each player and no additional statistics of performance,
- *Memoryless*, meaning that when the system updates players' ratings it considers only the game that was just played and not the history of previous game outcomes, and
- *Zero-sum*, meaning that the number of rating points gained by the winner of a game is always equal to the number of rating points lost by the loser.

To be clear, many rating systems used in theory or practice do *not* satisfy these properties for various reasons. Typically, this is because they intentionally augment a “base” rating system (which often does satisfy these three properties) to consider additional player statistics, game history, etc., in a way that trades simplicity and interpretability for improved empirical predictive accuracy of game outcomes [12, 5, 17, 6, 7, 19, 18]. Thus, this paper aims to study the space of “base” rating systems on which these more complex but practically accurate rating systems are built.

Additionally, a few of our results will consider the following property of rating systems:

- *Translation-invariant*, meaning that the win probabilities for two players are determined entirely by the additive difference between their ratings. (See Definition 27 for a formal definition.)

Translation invariance is common in practice; it implies that rating differences carry the same meaning for low-rated players as for high-rated players, and thus it makes the system more interpretable. Many systems specifically follow the convention that a 200-point rating difference maps to 0.75 win probability for the higher-rated player; this rule was first applied by the rating system ELO [9].

The goal of this paper is instead to continue a recent line of work on *incentive-compatibility* as a property of rating systems. The idea is that players experience rating as an incentive, and they will sometimes take strategic action to maximize their rating in the system. For example, Glicko-2 was recently subject to an attack called *volatility hacking* in which players achieved long-run boosts in their rating by deliberately losing certain games [8]. Some modern rating systems like URS address the possibility of strategic action by leaving parameters of their rating formula unpublished.² This has several disadvantages; for example, players are less able to predict their rating changes, which in turn makes it harder to detect transcription errors in game outcome reporting. Similarly, bugs in the implementation of a rating system may be harder to detect. It is desirable to achieve notions of incentive compatibility, even when the rating system in play is entirely public knowledge.

Our focus is on a new rating system property, informally stating that the system does not incentivize a rating-maximizing player to selectively seek games against opponents of one rating over another (even when the player currently feels over- or underrated by the system).

² Jeff Sonas, personal communication, 2022

One of our main conceptual contributions is just to *define* this property in a reasonable way: as we discuss shortly, the natural first attempt at a definition turns out to be too strong to be feasible. We then describe a different version of the definition, and in support of this definition, we show that there is a natural class of rating systems (most notably including *Sonas*, described shortly) that satisfy this alternate definition. In order to explain this progression in ideas formally, we next give details on our model of rating systems.

1.1.1 Formal Model of Rating Systems

We follow the common basic model that each player has a hidden “true rating,” as well as a visible “current rating” assigned by the system. Although current ratings naturally fluctuate over time, the goal of a rating system is to assign current ratings that are usually a reasonable approximation of true ratings. When two players play a game, their respective true ratings determine the probability that one beats the other, and their respective current ratings are used by the system as inputs to a function that determines the changes in the ratings of the winner and the loser.

Concretely, a rating system is composed of two ingredients. The first is the *skill curve* σ , which encodes the system’s model of the probability that a player of true rating x will beat a player of true rating y . The second is the *adjustment function* α , which takes the players’ current ratings as inputs, and encodes the number of points that the system awards to the first player and takes from the second player in the event that the first player beats the second player. We are not allowed to pair together an *arbitrary* skill curve and adjustment function; rather, a basic fairness axiom must be satisfied. This axiom is phrased in terms of an *expected gain function*, which computes the expected number of rating points gained by a player of current rating x and true rating x^* , when they play a game against a player of current rating y and true rating y^* . The fairness axiom states that, when any two correctly-rated players play a game, their expected rating change should be 0.

► **Definition 1 (Rating Systems).** A *rating system* is a pair (σ, α) , where:

- $\sigma : \mathbb{R}^2 \rightarrow [0, 1]$, the **skill curve**, is a continuous function that is weakly increasing in its first parameter and weakly decreasing in its second parameter. We assume the game has no draws, and thus σ must satisfy³

$$\sigma(x, y) + \sigma(y, x) = 1 \quad \text{for all } x, y. \text{ (}\sigma \text{ Draw-Free)}$$

- $\alpha : \mathbb{R}^2 \rightarrow \mathbb{R}_{\geq 0}$ is the **adjustment function**. We assume that every game has the possibility of changing the rating of one of the players, and thus α must satisfy

$$\alpha(x, y) + \alpha(y, x) > 0 \quad \text{for all } x, y. \text{ (Consequential)}$$

- Let $\gamma : \mathbb{R}^4 \rightarrow \mathbb{R}$ be the **expected gain function**, defined by

$$\gamma(x, x^* \mid y, y^*) := \alpha(x, y)\sigma(x^*, y^*) - \alpha(y, x)\sigma(y^*, x^*).$$

Then we require the following fairness axiom:

$$\gamma(x, x \mid y, y) = 0 \text{ for all } x, y.$$

³ We ignore the possibility of draws in this exposition purely for simplicity. In the case of draws, one could interpret a win as 1 victory point, a draw as 1/2 a victory point, and a loss as 0 victory points, and then interpret $\sigma(x, y)$ as the expected number of victory points scored by a player of true rating x against a player of true rating y . See [15] for more discussion, and for a rating systems that can handle even finer-grained outcomes than the win/loss/draw trichotomy.

The fairness axiom is only implicit in prior work. Other rating systems, including ELO, Glicko, etc., [12, 18, 14], often define their rating system by giving a skill curve and a *K-function*, which together imply an adjustment function in a sense we next explain. An adjustment function can be implicitly defined in this way if and only if it satisfies this fairness axiom. It will be technically convenient for us to be able to have both the explicit adjustment function α and the *K-function* to reference in our proofs.

1.1.2 *K*-Functions

Rearranging the fairness axiom, we get the identity

$$\sigma(x, y) = \frac{\alpha(y, x)}{\alpha(x, y) + \alpha(y, x)}.$$

Thus the adjustment function determines the skill curve (in the sense that, for a given adjustment function α , there is a unique skill curve σ that satisfies the fairness axiom). However, when designing a rating system, it is more intuitive to model the game with a skill curve first and then pick an adjustment function second. The skill curve does *not* fully determine the adjustment function; rather, the value of $\sigma(x, y)$ only determines the *ratio* between $\alpha(x, y)$ and $\alpha(y, x)$. The *K-function* is an auxiliary function that determines the scaling. We define it as the denominator part of the previous identity.⁴

► **Definition 2** (*K*-Functions). *For a rating system (σ, α) , its associated **K-Function** is defined by $K(x, y) := \alpha(x, y) + \alpha(y, x)$.*

Note that being consequential means that $K(x, y) > 0$ for all x, y .

Many rating systems in the literature do not *explicitly* define an adjustment function; rather, they define their adjustment function *implicitly* by instead giving the skill curve and the *K-function* [12, 9]. Most notably, Elo uses *K-factors*, which are a special case of our *K-functions*.⁵ See [21] for a discussion of the practical considerations behind choosing a good *K-function*.


1.2 The Elo and Sonas Rating Systems

The most famous and popular rating system used in practice is probably Elo. Its basis is a logistic skill curve: in a game between two players of true ratings x, y , Elo maps the quantity $x - y$ to a win probability for the first player using a logistic function. Many implementations of ELO, like the one used by the chess federation FIDE, also include a *thresholding rule*: e.g., if $x - y > 400$, then the FIDE implementation instead treats $x - y$ as exactly 400. In other words, it assumes the game is chaotic enough that no player ever has more than a 96% win probability over another. This thresholding also regularizes against extreme rating swings.

In 2002, Jeff Sonas published a famous critique of Elo [20], in which he analyzed a large sample of FIDE rated chess games among highly-rated players. He argued that the logistic curve is so flat within this player pool that a threshold-linear skill curve fits the data just as well, and hence is preferable due to its simplicity. We will henceforth call skill curves that are linear within some threshold *Sonas-like skill curves*:

⁴ One can also interpret *K-functions* through a gambling model where the players respectively put $\alpha(x, y), \alpha(y, x)$ of their rating points into a pot, and then the winner takes all the rating points in the pot. The *K-function* $K(x, y) = K(y, x)$ is the pot size.

⁵ We have chosen the name *K-function* to indicate that they generalize *K-factors*. *K* does not stand for anything; is an arbitrary variable name in the Elo adjustment function formula.



■ **Figure 1** Left: Elo skill curve (in yellow) overlaid with empirical data on the outcomes of FIDE-rated chess games. Right: Sonas skill curve (in red) overlaid with empirical data on the outcomes of FIDE-rated chess games in which both players were rated 2200 or higher. Figures by Jeff Sonas, taken from [21]. See [20, 22] for additional discussion.

► **Definition 3** (Sonas-like skill curves). *A skill curve is **Son_s-like** if there exist $a, s > 0$ such that for ratings $x, y \in \mathbb{R}$, when $|x - y| \leq s$, we have $\sigma(x, y) = ax - ay + 0.5$.*

Son_s's rating system specifically took $s = 400$, and assumes a flat skill curve outside this range; that is, $\sigma(x, y) = \sigma(y + 400, y)$ when $x > y + 400$. In 2011 [21], Son_s augmented his previous analysis with an interesting nuance: when *all* FIDE rated chess games are considered, including those where the participants have lower ratings, then the logistic curve is not so flat and Elo skill curves seem superior. Thus, the relative value of Elo and Son_s seems to depend a bit on the strength of the pool of players being modeled.

1.3 Strategyproof Rating Systems and Our Results

A system is strategyproof against a certain kind of undesirable strategic behavior if it does not incentivize a rating-maximizing player to perform that behavior. We next overview a recent case study, in which players in a real-world system discovered and implemented an attack on a rating system that was missing an important strategyproofness property.

1.3.1 Volatility Hacking

An obviously-desirable strategyproofness property is that a rating-maximizing player should always be incentivized to try their best to win each game they play. Interestingly, these incentives were recently discovered *not* to hold in the popular Glicko-2 rating system.

Glicko-2 [12] is a higher-dimensional version of Elo that tracks a *volatility parameter* for each player. A higher volatility parameter indicates that the system is more uncertain about the player's true rating, which in turn amplifies the adjustment function. This volatility parameter was repeatedly shown to improve predictive accuracy of systems, and so it has been widely adopted. But recently, an attack has been discovered called *volatility hacking*. A player first strategically loses many games, lowering their rating to the point where they only face much lower-rated opponents and can essentially win or lose at will. Then, the player alternates wins and losses, which boosts their volatility parameter as much as they like while keeping their rating roughly stable. Finally, they win a few games, and due to their enormous volatility parameter their rating skyrockets well above its initial value.

Volatility hacking was used to attack the game Pokémon Go [25], which uses a version of Glicko-2 to rank players. An interesting recent paper by Ebtekar and Liu [8] shows that some other popular rating systems are also vulnerable to volatility hacking attacks, and it proposes concrete fixes to the handling of volatility parameters that would make a Glicko-2-like system (with a volatility parameter) strategyproof against volatility hacking.

1.3.2 Opponent Indifference and First Main Result

This paper is about securing rating systems against a less pernicious but likely much more common type of attack that we will call *opponent selection*. Many matchmaking systems provide players with some control over their next opponent. For example, some online games use an interface in which players view a public list of active challenges from opponents of various ratings, and they can accept any challenge they like. Alternately, some games use random matchmaking, but then let players abort a game without penalty after glimpsing their opponent's rating. If players have any control over their next opponent's rating, then it is clearly undesirable for a rating system to allow a player to boost their rating by exercising this control strategically. For example, in the Elo rating system, one can compute that the magnitude of expected gain is greatest when a (correctly-rated) opponent's true rating is exactly halfway between a player's current rating and their true rating.

Which rating systems are *immune* to opponent selection incentives? The natural definition would capture systems where the a player's expected gain doesn't depend on their opponent's rating, so long as that opponent is correctly rated. Formally:

► **Definition 4 (Opponent Indifference).** A rating system (σ, α) with expected gain function γ is **opponent indifferent** if there is a function $\gamma^* : \mathbb{R}^2 \rightarrow \mathbb{R}$ such that

$$\gamma(x, x^* \mid y, y) = \gamma^*(x, x^*) \text{ for all } x, x^*, y.$$

Before we proceed, let us be more specific about the kind of opponent selection attacks that become impossible in a rating system that satisfies this definition. The premise is that the players in a system are not usually correctly rated, i.e., their current rating differs significantly from their true rating. One reason is due to random walk effects, where a player's current rating will naturally drift over time. Another reason is that a player's true rating might change from day to day due to external factors like rest, stress, distraction, etc.⁶ An opponent selection attack would have a player seek opponents of one rating when they feel overrated by the system, and another when they feel underrated, causing their rating fluctuations to hold a significantly higher average than they would without strategic behavior.

Having motivated the axiom of opponent indifference, our next goal is to check whether it is satisfied by any actual rating systems, and whether these systems are “reasonable.” Although opponent indifferent rating systems exist (this follows from Lemma 24 in this paper), we actually argue that they are *not* sufficiently expressive to be interesting. Our reasoning is as follows. Most interesting games exhibit *skill chains*. In chess, for example, a hobbyist can beat a beginner (say) 95% of the time, and an expert can beat a hobbyist 95% of the time, and a master can beat an expert 95% of the time, and so on. Most rating systems allow, in theory, for an infinite chain of players of ascending ratings, with each player beating the next with reasonably high probability. We can formalize a version of this property as follows:

► **Definition 5 (Full Scale).** A rating system (σ, α) has **full scale** if there exists $p > 0.5$ and an infinite ascending chain of ratings $r_1 < r_2 < \dots$ such that $\sigma(r_i, r_{i-1}) \geq p$ for all i .

Our critique of opponent indifference is that it is incompatible with full scale, and thus cannot really capture games like chess.

⁶ Experts have estimated that chess strength fluctuates by about ± 200 Elo rating points (in the FIDE system) on a given day due to external factors [4]. However, we have not been able to find statistical validation of these estimates, so they should be taken as informal.

► **Theorem 6** (First Main Result – See Theorem 25 in the body). *There is no rating system that simultaneously satisfies opponent indifference and full scale.*

Although we note that every major rating system we are aware of satisfies full scale, one could object at this point that full scale is too strong. One might not expect an *infinite* chain of ratings $r_1 < r_2 < \dots$ as in the definition of full scale to exist, but rather only a “reasonably long” chain of ratings. The argument in our impossibility theorem (Theorem 25) arguably restricts these rating chains to be unreasonably short. For example: if we take parameter $p = 0.75$, i.e., we want a chain of ratings $r_1 < r_2 < \dots$ where each rating beats the previous one with probability ≥ 0.75 , this chain can have only length 3. This implies that an opponent indifferent rating system following the ELO convention where $\sigma(x + 400, x) = 0.75$ would have a maximum possible rating spread of only 600 points.

1.3.3 P Opponent Indifference and Second Main Result

In light of Theorem 6, our goal is now to weaken the definition of opponent indifference in a way that escapes the impossibility result of Theorem 6, while still providing an effective type of immunity against opponent selection attacks. Fortunately, we argue that a reasonable relaxation of opponent indifferent is already implicit in the *thresholding effects* used in implementations of Elo and in Sonas. These systems hardcode a rating difference at which the rating system punts: the system declares that chaotic effects dominate, and ratings are no longer a good predictor of outcome probabilities. A natural relaxation is to enforce opponent indifference *only* between opponents in the “non-chaotic” rating regime.

We call this weaker version P *opponent indifference*, where the parameter P controls the threshold beyond which we no longer require opponent indifference to hold. In the following, let us say that ratings x, y are P -close if we have $\sigma(x, y) \in (0.5 - P, 0.5 + P)$. Formally:⁷

► **Definition 7** (P Opponent Indifference). *For a parameter $P \in (0, 0.5]$, a rating system is P **opponent indifferent** if there is a function $\gamma^* : \mathbb{R}^2 \rightarrow \mathbb{R}$ such that*

$$\gamma(x, x^* \mid y, y) = \gamma^*(x, x^*)$$

for any x, x^*, y with x, y and x^*, y both P -close.

P opponent indifference might feel like a light tweak on opponent indifference; it places the exact same requirements as opponent indifference on all “reasonable” games that might be played (according to a thresholding rule). Thus it provides immunity against opponent selection attacks in the cases of interest. But, perhaps surprisingly, this relaxation is enough to escape impossibility. We prove the following characterization theorem.

► **Definition 8** (P Separable). *A skill curve σ is P **separable** if there is a weakly increasing function $\beta : \mathbb{R} \rightarrow \mathbb{R}$ such that $\sigma(x, y) = \beta(x) - \beta(y) + 0.5$ for all P -close x, y .*

► **Definition 9** (P Constant). *For $P \in (0, 0.5]$, a K -function is P **constant** if for all x, y such that $\sigma(x, y) \in (0.5 - P, 0.5 + P)$, $K(x, y) = C$ for some constant C .*

► **Theorem 10** (Second Main Result – See Theorem 32 in the body). *A nontrivial rating system (σ, α) is P opponent indifferent if and only if σ is P separable and its K -function is P constant.*

⁷ See Definition 31 for the equivalent technical definition used in the paper.

It is relatively easy from this characterization theorem to show that P opponent rating systems exist, and that they can exhibit full scale (even with respect to any given parameter P).⁸ In the next part, we show that these desirable properties continue to hold even under a natural *strengthening* of P opponent indifference.

1.3.4 Strong P Opponent Indifference and Third Main Result

P opponent indifference requires indifference between any two opponents within threshold, *assuming those opponents are correctly rated*. One might want to more strictly require indifference between two opponents that are both *incorrectly rated by the same amount*. Formally:

► **Definition 11** (Strong P Opponent Indifference). *For a parameter P , a rating system is strongly P opponent indifferent if there is a function $\gamma^* : \mathbb{R}^3 \rightarrow \mathbb{R}$ such that*

$$\gamma(x, x^* \mid y, y + \delta) = \gamma^*(x, x^*, \delta)$$

for all $x, x^*, y, y + \delta$ with x, y and $x^*, y + \delta$ both P -close.

Note that P opponent indifference is the special case where $\delta = 0$. Fortunately, this strengthening remains tenable, as shown by the following theorem:

► **Theorem 12** (Third Main Result – See Corollary 41). *For any $0 < P \leq 0.5$ and any rating system (σ, α) , the rating system is strongly P opponent indifferent if and only if it is Sonas-like with a P constant K -function.*

Conceptually, we interpret this theorem as a significant technical point in favor of Sonas. For games where Elo and Sonas are equally preferable in terms of accurately modeling the player pool, the Sonas model has the additional advantage of fighting opponent selection attacks, which may give a reason why it should be favored.

In Corollary 30, we also discuss the analogous strengthening of (general) opponent indifference, and we strengthen our impossibility result to show that no nontrivial strong (general) opponent indifferent rating systems exist, whether or not they satisfy full scale.

1.4 Further Details on Previous Work

Here we supply a bit more technical detail on some prior work referenced in the above discussion.

1.5 Applications

The most direct and common use of rating systems is in *matchmaking*, where a system tracks estimates of player strength in order to create well-balanced and engaging matches. Some recent work has focused on *team* matchmaking, updating estimates of player strength when the system must create two opposing *groups* worth of players. See [3] and references within for examples.

Rating systems have also been applied in economic contexts, such as pricing algorithms [26]. Agents in the economy are given ratings, and items they might want to purchase are also given ratings. An item is offered to an agent, at a price determined by the rating of

⁸ This follows e.g. as a consequence of our Corollary 41, which implies that the Sonas rating system itself is P opponent indifferent (and it has full scale); we discuss this more next.

the item. If the agent buys the item, then the item “wins” the interaction; its rating (i.e., the system’s appraisal of its value) increases, while the agent’s rating (i.e., the system’s estimate of their stinginess) decreases. Alternately, if the agent refuses the item, then the agent “wins” the interaction, and the item’s rating decreases. A related use case is item response theory [11], where the goal is to categorize the difficulty of tasks, e.g., answering various questions on a test. Here we again view each agent answering each question as a game, where the agent “wins” the interaction if they get the question right, and the question “wins” the interaction if the agent gets it wrong. The system then updates its beliefs about the agent’s capabilities or the question’s difficulty accordingly.

1.6 Related Strategyproof Rating Mechanisms

Closely related to *rating systems* are *ranking systems*, where the goal is simply to order agents by relative skill (construed broadly; e.g., this would include rankings of colleges by “quality”). Although ranking systems often work by assigning agents a numeric score that could be interpreted as a rating, the main difference from our work is that they are *offline*: they post-process data into a ranking, rather than updating ratings in an *online* fashion in response to newly arriving data. Agents who are aware of a ranking algorithm may act strategically to optimize their position, and thus, incentive-compatibility is an important feature of these algorithms. For some work in this area, see [16] and references therein.

In our work, we do not assume any prior knowledge of agent skill levels, and we do not necessarily assume that agents even know their own skill level. If we assume that system or agents *do* have some prior knowledge of ratings, we might want a mechanism that incentivizes them to report their beliefs truthfully. There is a related line of work developing such mechanisms [10, 23, 24].

1.7 Open Problems and Future Directions

We conclude our introduction by suggesting a few open problems arising from our work.

1. What systems can satisfy opponent indifference properties if we relax our model to allow rating systems that are not one-dimensional, memoryless, or zero-sum?
2. Can opponent selection attacks be avoided by systems that separately track a *display rating*, which is visible to players and which they may try to maximize, from a *matchmaking rating*, used internally by the system to adjust ratings? How different would the display ratings and matchmaking ratings need to be?⁹
3. Can opponent indifference properties be achieved simultaneously with other desirable properties of rating systems? Other natural properties to consider might include fast convergence of current ratings to true ratings, or other notions of strategyproofness, like the one considered by Ebtekar and Liu [8].

2 Preliminaries

Some of our results in the introduction reference *nontrivial* rating systems. We define these formally:

► **Definition 13** (Trivial and nontrivial). *A rating system is **nontrivial** if for all x there exists a y such that $\sigma(x, y) = 0.5$, or **nontrivial** otherwise.*

⁹ We are grateful to an anonymous reviewer who suggested this open problem.

21:10 Opponent Indifference in Rating Systems

In the introduction, we discuss versions of properties that only hold over a restricted domain of ratings (like P opponent indifference, P constant, etc). In the technical part of this paper, we will need to similarly generalize other functions, and so the following (slightly informal) language will be helpful:

► **Definition 14** (Property over (A, B)). *Given a property defined over a set of inputs in \mathbb{R} , we say that a property holds **over** (A, B) if it holds whenever the inputs are taken in the interval (A, B) . We will equivalently say that (A, B) is a property interval for the rating system.*

For example, a rating system is nontrivial over (A, B) if there exist $x, y \in (A, B)$ with $\sigma(x, y) \neq 0.5$, and in this case we say that (A, B) is a nontrivial interval for the rating system.

The definition of the expected gain function in the introduction is phrased in a way that makes its intuitive meaning clear. However, the following alternate characterization in terms of the K -function will be useful in some of our proofs.

► **Lemma 15.** *The expected gain function for a rating system (σ, α) satisfies*

$$\gamma(x, x^* \mid y, y^*) = K(x, y)(\sigma(x^*, y^*) - \sigma(x, y))$$

Proof. We compute:

$$\begin{aligned} \gamma(x, x^* \mid y, y^*) &= \alpha(x, y)\sigma(x^*, y^*) - \alpha(y, x)\sigma(y^*, x^*) \\ &= K(x, y)\sigma(y, x)\sigma(x^*, y^*) - K(y, x)\sigma(x, y)\sigma(y^*, x^*) && \text{Def of } \alpha \text{ using } K\text{-function} \\ &= K(x, y)(\sigma(y, x)\sigma(x^*, y^*) - \sigma(x, y)\sigma(y^*, x^*)) && K \text{ is symmetric} \\ &= K(x, y)((1 - \sigma(x, y))\sigma(x^*, y^*) - \sigma(x, y)(1 - \sigma(x^*, y^*))) && \sigma \text{ Draw-Free} \\ &= K(x, y)(\sigma(x^*, y^*) - \sigma(x, y)\sigma(x^*, y^*) - \sigma(x, y) + \sigma(x, y)\sigma(x^*, y^*)) \\ &= K(x, y)(\sigma(x^*, y^*) - \sigma(x, y)). \quad \blacktriangleleft \end{aligned}$$

3 Opponent Indifference

In this section we will characterize all opponent indifferent rating systems. Some of our lemmas are proved with extra generality “over (A, B) ”; this generality will become useful in the following section.

3.1 Opponent Indifference vs. Full Scale

This section is about the incompatibility between the opponent indifference and full scale axioms discussed in the introduction. Our main technical lemma in this section is Lemma 19, which shows that any nontrivial opponent indifferent rating system has a constant K -function; this is proved by analyzing the expected gain function under the restriction implied by opponent indifference. We then use this result about the K -function to reveal the structure of the skill curve. This ultimately makes full-scale impossible (Theorem 25).

► **Lemma 16.** *If a rating system is both nontrivial and opponent indifferent over (A, B) where $\sigma(x, z) = 0.5$ and $\sigma(y, z) = 0.5$, then $\sigma(x, y) = 0.5$ for $x, y, z \in (A, B)$.*

Proof. Opponent indifference means that the expected gain only depends on the first two inputs, so

$$\begin{aligned} \gamma(x, y \mid z, z) &= \gamma(x, y \mid x, x) && \text{Opponent Indifference} \\ K(x, z)(\sigma(y, z) - \sigma(x, z)) &= K(x, x)(\sigma(y, x) - \sigma(x, x)) && \text{Lemma 15} \\ K(z, x)(0.5 - 0.5) &= K(x, x)(\sigma(y, x) - 0.5) && \sigma \text{ Draw-Free} \\ 0 &= K(x, x)(\sigma(y, x) - 0.5) \end{aligned}$$

Since the K -function is always positive as rating systems are consequential, we must have $\sigma(x, y) = 0.5$. \blacktriangleleft

► **Lemma 17.** *If a rating system is both nontrivial and opponent indifferent over (A, B) , then for all $x \in (A, B)$, there exists $y \in (A, B)$ such that $\sigma(x, y) \neq 0.5$.*

Proof. Suppose there was a rating $x \in (A, B)$ such that $\sigma(x, y) = 0.5$ for all $y \in (A, B)$. There exists $z, z^* \in (A, B)$ such that $\sigma(z, z^*) \neq 0.5$ by nontriviality. We then have $\sigma(z, x) = 0.5$ and $\sigma(z^*, x) = 0.5$. This implies $\sigma(z^*, z) = 0.5$ by Lemma 16 which is a contradiction. \blacktriangleleft

► **Lemma 18.** *If a rating system is both nontrivial and opponent indifferent over (A, B) , then for $x, y \in (A, B)$ such that $\sigma(x, y) \neq 0.5$, $K(x, x) = K(x, y)$.*

Proof. Let $x, y \in (A, B)$ where $\sigma(x, y) \neq 0.5$. From the definition of opponent indifference, the expected gain does not depend on its latter two parameters. We therefore have

$$\begin{aligned} \gamma(x, y \mid x, x) &= \gamma(x, y \mid y, y) && \text{Opponent Indifference} \\ K(x, x)(\sigma(y, x) - \sigma(x, x)) &= K(x, y)(\sigma(y, y) - \sigma(x, y)) && \text{Lemma 15} \\ K(x, x)(\sigma(y, x) - 0.5) &= K(x, y)(0.5 - \sigma(x, y)) \\ K(x, x)(1 - \sigma(x, y) - 0.5) &= K(x, y)(0.5 - \sigma(x, y)) && \sigma \text{ Draw-Free} \\ K(x, x)(0.5 - \sigma(x, y)) &= K(x, y)(0.5 - \sigma(x, y)) \\ K(x, x) &= K(x, y). && \blacktriangleleft \end{aligned}$$

► **Lemma 19.** *If a rating system (σ, α) is opponent indifferent over (A, B) and nontrivial over (A, B) , then $K(x, y)$ is constant over (A, B) .*

Proof. Let $x, y \in (A, B)$. We want to show that $K(x, x) = K(x, y)$. This means that the K -function depends only on its first input. Since additionally K is symmetric in its two parameters, it must in fact be constant over (A, B) . Lemma 18 proves the statement in the case that $\sigma(x, y) \neq 0.5$, so we need only cover the case that $\sigma(x, y) = 0.5$.

Lemma 17 shows that there exists $z \in (A, B)$ such that $\sigma(x, z) \neq 0.5$. Lemma 16 implies $\sigma(z, y) \neq 0.5$ by contrapositive as $\sigma(x, z) \neq 0.5$ and $\sigma(x, y) = 0.5$. We can then evaluate

$$\begin{aligned} \gamma(x, z \mid y, y) &= \gamma(x, z \mid z, z) \\ K(x, y)(\sigma(z, y) - \sigma(x, y)) &= K(x, z)(\sigma(z, z) - \sigma(x, z)) \\ K(x, y)(\sigma(z, y) - 0.5) &= K(x, z)(0.5 - \sigma(x, z)) \end{aligned}$$

and

$$\begin{aligned} \gamma(y, z \mid z, z) &= \gamma(y, z \mid x, x) \\ K(y, z)(\sigma(z, z) - \sigma(y, z)) &= K(y, x)(\sigma(z, x) - \sigma(y, x)) \\ K(y, z)(0.5 - \sigma(y, z)) &= K(y, x)(\sigma(z, x) - 0.5) \\ K(y, z)(\sigma(z, y) - 0.5) &= K(y, x)(0.5 - \sigma(x, z)) \end{aligned}$$

21:12 Opponent Indifference in Rating Systems

Since the K -function is symmetric, we let $a = K(x, y) = K(y, x)$ and by Lemma 18, we let $b = K(x, z) = K(y, z)$. We then let $c = (\sigma(z, y) - 0.5)$ and $d = (0.5 - \sigma(x, z))$, giving us

$$a \cdot c = b \cdot d$$

and

$$b \cdot c = a \cdot d$$

The previous two equations imply that $a = b$, and thus $K(x, y) = K(x, z) = K(x, x)$ by Lemma 18. ◀

We comment that the nontriviality hypothesis in the previous lemmas is indeed necessary, since otherwise $K(x, y)$ can be essentially any symmetric function over (A, B) and the expected gain will still be 0.

► **Lemma 20.** *If a rating system (σ, α) is opponent indifferent over (A, B) , then*

$$\sigma(x, y) = \sigma(x, z) - \sigma(y, z) + 0.5$$

for all $x, y, z \in (A, B)$.

Proof. If the system is trivial over (A, B) then the lemma holds immediately, so assume nontriviality. From the definition of opponent indifference the expected gain only depends on a player's current rating and true rating, x and x^* , so long as all ratings are within the interval (A, B) . So we have:

$$\begin{aligned} \gamma(x, x^* | y, y) &= \gamma(x, x^* | z, z) && \text{Opponent Indifference} \\ K(x, y)(\sigma(x^*, y) - \sigma(x, y)) &= K(x, z)(\sigma(x^*, z) - \sigma(x, z)) && \text{Lemma 15} \\ C \cdot (\sigma(x^*, y) - \sigma(x, y)) &= C \cdot (\sigma(x^*, z) - \sigma(x, z)) && \text{Lemma 19.} \end{aligned}$$

Now consider the possible setting $x^* = y$. Under this, we continue:

$$\begin{aligned} C \cdot (\sigma(y, y) - \sigma(x, y)) &= C \cdot (\sigma(y, z) - \sigma(x, z)) \\ 0.5 - \sigma(x, y) &= \sigma(y, z) - \sigma(x, z) \\ \sigma(x, y) &= \sigma(x, z) - \sigma(y, z) + 0.5. \end{aligned}$$

► **Definition 21 (Separable).** *A skill curve is **separable** if there is a function $\beta : \mathbb{R} \rightarrow [C, C + 0.5]$ for some constant C such that*

$$\sigma(x, y) = \beta(x) - \beta(y) + 0.5.$$

for all x, y . β is called a **bisector** of the skill curve.

Note that bisectors are not unique: if β is a bisector for σ , then any vertical translation of β is also a bisector. However, all bisectors differ by this translation:

► **Lemma 22.** *For any skill curve σ that is separable over (A, B) , if β, β' are both bisectors of σ over (A, B) , then there is a constant C such that $\beta'(x) = \beta(x) + C$ for all $x \in (A, B)$.*

Proof. Fix an arbitrary $y \in (A, B)$. We have

$$\sigma(x, y) = \beta'(x) - \beta'(y) + 0.5 = \beta(x) - \beta(y) + 0.5$$

and so, rearranging, we get

$$\beta'(x) = \beta(x) + (\beta'(y) - \beta(y)).$$

Now the claim follows by taking $C := \beta'(y) - \beta(y)$. ◀

► **Lemma 23.** *If a rating system (σ, α) is opponent indifferent over (A, B) , then σ is separable over (A, B) and all bisectors are continuous.*

Proof. Suppose (σ, α) is opponent indifferent over (A, B) . Let $m \in (A, B)$ be an arbitrary constant. From Lemma 20 we have

$$\sigma(x, y) = \sigma(x, m) - \sigma(y, m) + 0.5.$$

Therefore $\beta(x) := \sigma(x, m)$ is a bisector for σ over (A, B) , and since σ is continuous, β is continuous as well. Finally, Lemma 22 implies that since one bisector is continuous, all bisectors are continuous. ◀

► **Lemma 24.** *If a rating system's K -function is constant over (A, B) and the skill curve is separable over (A, B) then the rating system is opponent indifferent over (A, B) .*

Proof. Plugging the equations into the expected gain function for $x, x^*, y \in (A, B)$, we have

$$\begin{aligned} \gamma(x, x^* \mid y, y) &= K(x, y)(\sigma(x^*, y) - \sigma(x, y)) && \text{Lemma 15} \\ &= C \cdot (\beta(x^*) - \beta(y) + 0.5 - (\beta(x) - \beta(y) + 0.5)) \\ &= C \cdot (\beta(x^*) - \beta(x)). \end{aligned}$$

Thus γ depends only on its first two parameters x and x^* , implying opponent indifference. ◀

Note that Lemma 24 implies that nontrivial opponent indifferent rating systems do indeed exist, as we can choose a constant K -function and suitable bisector function β over the entire interval $(-\infty, \infty)$. However, as discussed in the introduction, our next theorem shows that no opponent indifferent rating system can exhibit full scale:

► **Theorem 25.** *No opponent indifferent rating system (over $(-\infty, \infty)$) has full scale.*

Proof. Let (σ, α) be an opponent indifferent rating system and fix some $p \in (0.5, 1)$. Our plan is to prove that for some integer N depending only on p , there does not exist a chain of ratings $r_1 < \dots < r_N$ such that for all $1 < i \leq N$, we have $\sigma(r_i, r_{i-1}) = p$. More specifically, our strategy is to prove that

$$\sigma(r_N, r_1) \geq (N-1)p - \frac{N-2}{2}. \tag{1}$$

Since $\sigma(r_N, r_1) \leq 1$, this implies

$$\begin{aligned} (N-1)p - \frac{N-2}{2} &\leq 1 \\ Np - p - \frac{N}{2} + 1 &\leq 1 \\ N \left(p - \frac{1}{2} \right) &\leq p \\ N &\leq \frac{2p}{2p-1} \end{aligned}$$

which is an upper bound for N depending only on p , as desired. It now remains to prove equation (1). We do so by induction on N :

■ (Base Case, $N = 1$) We have $\sigma(r_1, r_1) = 1/2$, as desired.

21:14 Opponent Indifference in Rating Systems

- (Inductive Step) By Lemma 23 the skill curve is separable, and so for some bisector β we have

$$\begin{aligned}
 \sigma(r_{N+1}, r_1) &= \sigma(r_{N+1}, r_N) - \sigma(r_1, r_N) + 0.5 && \text{Lemma 20} \\
 &\geq p - (1 - \sigma(r_N, r_1)) + 0.5 && \sigma \text{ Draw-Free} \\
 &\geq p - \left(1 - \left((N-1)p - \frac{N-2}{2}\right)\right) + 0.5 && \text{Inductive Hypothesis} \\
 &= p - 1 + (N-1)p - \frac{N-2}{2} + 0.5 \\
 &= Np - \frac{N-1}{2}
 \end{aligned}$$

which verifies (1) for $N+1$. ◀

We note the setting $N \leq \frac{2p}{2p-1}$ that arises in this theorem, which for a given p controls the maximum possible length of a skill chain that may exist in an opponent indifferent system. In particular, plugging in $p = 0.75$, we get $N \leq \frac{1.5}{0.5} = 3$.

3.2 Impossibility of Strong Opponent Indifference

Here, we consider an even stronger notion of opponent indifference. While opponent indifference enforces that we are indifferent between two *correctly-rated* opponents, we could more strongly require indifference between two opponents *misrated by the same amount*. In particular, the following definition enforces indifference between opponents who are overrated by δ :

► **Definition 26** (Strong Opponent Indifference). *A rating system is **strongly opponent indifferent** if there is a function $\gamma^* : \mathbb{R}^3 \rightarrow \mathbb{R}$ such that, for all $x, x^*, y, y + \delta$, we have*

$$\gamma(x, x^* \mid y, y + \delta) = \gamma^*(x, x^*, \delta).$$

A strongly opponent indifferent rating system is also opponent indifferent (by considering $\delta = 0$), and thus by Theorem 25, it cannot exhibit full scale. We will prove an even stronger impossibility theorem. While this is a slight detour (as we have arguably already shown that strong opponent indifference is undesirable), the results in this section are used in the following section. As before, we prove our intermediate lemmas over arbitrary intervals (A, B) , which provides flexibility that will be helpful later in the paper.

► **Definition 27** (Translation Invariant). *A rating system (σ, α) is **translation invariant** if there is a function $\sigma^* : \mathbb{R} \rightarrow [0, 1]$ with $\sigma(x, y) = \sigma^*(x - y)$ for all x, y .*

► **Theorem 28**. *If a rating system is strongly opponent indifferent over (A, B) , then it is translation invariant over (A, B) .*

Proof. If the system is trivial over (A, B) then the lemma is immediate, so assume non-triviality over (A, B) . Let (σ, α) be a strongly opponent indifferent rating system. To show translation invariance over (A, B) , let x, x^*, δ be such that $x, x^*, x + \delta, x^* + \delta \in (A, B)$. Since strong opponent indifference implies opponent indifference, Lemmas 15 and 19 apply. We can therefore compute:

$$\begin{aligned}
 \gamma(x, x^* \mid y, y + \delta) &= K(x, y)(\sigma(x^*, y + \delta) - \sigma(x, y)) && \text{Lemma 15} \\
 &= C \cdot (\sigma(x^*, y + \delta) - \sigma(x, y)) && \text{Lemma 19.}
 \end{aligned}$$

From the definition of strong opponent indifference, the expected gain function depends only on the first two parameters and the difference between its latter two parameters, and so we have

$$\begin{aligned}
\gamma(x, x^* \mid x^*, x^* + \delta) &= \gamma(x, x^* \mid x, x + \delta) \\
C \cdot (\sigma(x^*, x^* + \delta) - \sigma(x, x^*)) &= C \cdot (\sigma(x^*, x + \delta) - \sigma(x, x)) && \text{previous equation} \\
\sigma(x^*, x^* + \delta) - \sigma(x, x^*) &= \sigma(x^*, x + \delta) - 0.5 \\
\sigma(x, x^*) &= \sigma(x, x^* + \delta) - \sigma(x^*, x + \delta) + 0.5 \\
\sigma(x, x + \delta) - \sigma(x^*, x + \delta) + 0.5 &= \sigma(x^*, x^* + \delta) - \sigma(x^*, x + \delta) + 0.5 && \text{Lemma 20} \\
\sigma(x, x + \delta) &= \sigma(x^*, x^* + \delta),
\end{aligned}$$

which implies translation invariance over (A, B) . ◀

► **Theorem 29.** *If a rating system (σ, α) is strongly opponent indifferent over (A, B) , then σ has a bisector that is linear over (A, B) (i.e., on the interval (A, B) it coincides with a function of the form $\beta(x) = mx$ for some constant m).*

Proof. By considering an appropriate horizontal translation of our rating system, we may assume without loss of generality that our opponent indifferent interval is symmetric about the origin, i.e., it has the form $(-A, A)$. Let β be an arbitrary function that bisects σ over $(-A, A)$, and further assume without loss of generality that β intersects the origin, i.e., $\beta(0) = 0$. Our goal is to prove that, for all $x, y \in (-A, A)$, we have $\beta(x + y) = \beta(x) + \beta(y)$. Since we already have that β is continuous, this implies linearity of β over $(-A, A)$ (see Lemma 42 in Appendix A for details).

Let $x, y, x + y \in (-A, A)$. We then have

$$\begin{aligned}
\sigma(x + y, x) &= f(y) && \text{Theorem 28} \\
\beta(x + y) - \beta(x) &= f(y) && \text{Lemma 23} \\
\beta(x + y) &= \beta(x) + f(y). && (2)
\end{aligned}$$

Additionally, by reversing the roles of x and y , we have

$$\beta(x + y) = f(x) + \beta(y). \quad (3)$$

Thus, combining (2) and (3), we get

$$f(x) + \beta(y) = \beta(x) + f(y).$$

It follows from equation (2) that $f(0) = 0$, and we recall that $\beta(0) = 0$. Therefore, plugging $y = 0$ into the previous equation, we get

$$\begin{aligned}
f(x) + \beta(0) &= \beta(x) + f(0) \\
f(x) &= \beta(x)
\end{aligned}$$

and so, recombining with (2), we have $\beta(x + y) = \beta(x) + \beta(y)$, as desired. ◀

► **Corollary 30.** *A strongly opponent indifferent rating system (over $(-\infty, \infty)$) must be trivial.*

Proof. If a rating system is strongly opponent indifferent, then by Theorem 29 there exists a bisector β that is linear over $(-\infty, \infty)$. Additionally, the range of β is contained within $[C, C + 0.5]$ for some constant C . The only such functions are those with slope 0. Thus

$$\sigma(x, y) = 0 - 0 + 0.5 = 0.5$$

for all $x, y \in \mathbb{R}$, and so it is trivial. ◀

4 P Opponent Indifference

Recall that P opponent indifference is a relaxation of (full) opponent indifference, in which we only require indifference among opponents who are reasonably evenly matched.

► **Definition 31** (P Opponent Indifference). For $P \in (0, 0.5]$, a rating system is **P opponent indifferent** if the rating system is opponent indifferent over (A, B) for all $A < B$ where¹⁰

$$\sigma(A, B) > 0.5 - P.$$

4.1 Characterization of P Opponent Indifference

Our next goal is to prove the following characterization theorem:

► **Theorem 32** (Characterization of P Opponent Indifferent Rating Systems). A nontrivial rating system is P opponent indifferent if and only if the skill curve is P separable and the K -function is P constant.

We work towards a proof with some intermediate structural lemmas.

► **Lemma 33.** If a rating system (σ, α) is P opponent indifferent then every real number is within an opponent indifferent interval.

Proof. Let $b \in \mathbb{R}$. Since σ is continuous, there exists $0 < \epsilon$ such that $0.5 \geq \sigma(b - \epsilon, b + \epsilon) > 0.5 - P$. Thus $(b - \epsilon, b + \epsilon)$ is an opponent indifferent interval. ◀

► **Lemma 34.** If a rating system (σ, α) is P opponent indifferent, then σ is P separable with a continuous bisector.

Let us quickly discuss this lemma statement before we begin its proof. By Lemma 23, for each P -interval (A, B) , there exists a continuous function β that bisects σ on this interval. But this lemma is claiming something stronger: that there is a *single* continuous function β that bisects σ on *all* P -intervals simultaneously. This strengthening requires a bit of topology to prove, but is overall fairly straightforward.

Proof of Lemma 34. We construct a bisector β as follows. First, we arbitrarily fix a point; say, $\beta(0) = 0$, and all other values of β are currently undefined. Then, iterate the following process. Let x be the current supremum of the points on which β has been defined. Let $\epsilon > 0$ be the largest value such that $(x - \epsilon, x + \epsilon)$ is a P -interval. Using Lemma 23, there exists an infinite family of bisectors over $(x - \epsilon, x + \epsilon)$, which are vertical translations of each other. However, since $(x - \epsilon, x + \epsilon)$ intersects at least one previously-defined point, only one bisector in this family is consistent with the previously-selected values of β . We may extend β to the interval $(x - \epsilon, x + \epsilon)$ using this particular bisector, and repeat infinitely.

We claim that, for every nonnegative number r , there is a constant c_r such that the value of $\beta(r)$ is defined after finitely many iterations of this process. To see this, let $D \subseteq \mathbb{R}$ be the set of points with this property, and suppose for contradiction that the supremum of D is a finite real number r^* . Note that D is the union of open intervals, and thus D is open, which means it does not contain its supremum. Since $r^* \notin D$, there is no constant c_{r^*} for

¹⁰Technically, this definition differs slightly from Definition 7 of partial opponent indifference in the introduction, because it forces all pairs among x, x^*, y to be P -close, whereas Definition 7 only explicitly forces x, y , and x^*, y to be P -close. However, it follows as an easy corollary of the results in this section that the two definitions are equivalent (see Lemma 43 in Appendix A for details).

which $\beta(r^*)$ is defined after c_{r^*} iterations. Additionally, since r^* is the supremum of D , any interval of the form $(r^* - \varepsilon, r^* + \varepsilon)$ intersects D . Let $r \in (r^* - \varepsilon, r^* + \varepsilon) \cap D$. Then after $c_r + 1$ iterations, we would define $\beta(r^*)$. Since $c_r + 1$ is finite, this implies $r^* \in D$, which completes the contradiction.

By a symmetric process, we may then extend β to negative inputs. ◀

► **Lemma 35.** *For a nontrivial P opponent indifferent rating system, every opponent indifferent interval (A, B) has a constant K -function over (A, B) .*

Proof. The rating system is either nontrivial or trivial over (A, B) . If the rating system is nontrivial, the result is immediate by Lemma 19. Thus the remaining case is when the system is trivial over (A, B) , but nontrivial overall.

We have $\sigma(A, B) = 0.5$. Since σ is not identically 0.5, there exist $B' > B, A' < A$ such that $\sigma(A', B') < 0.5$. Moreover, since σ is continuous, we may specifically choose values A', B' satisfying

$$0.5 - P < \sigma(A', B') < 0.5.$$

So (A', B') is a P -interval, and σ is nontrivial over (A', B') . By Lemma 19, we thus have a constant K -function over (A', B') . Since $(A, B) \subseteq (A', B')$, the lemma follows. ◀

We are now ready to prove Theorem 32.

Proof of Theorem 32. (\longrightarrow) Lemma 34 implies that σ is P separable, so it only remains to show that the K -function is P constant. We prove this in two steps. First, let $f(x) = K(x, x)$, and we will prove that $f(x)$ is constant. Since every x lies in an opponent indifferent interval (Lemma 33), by Lemma 35 f is constant on that interval. In particular this implies f is differentiable at x , with $f'(x) = 0$, and thus f is constant. Finally: for any x, y in the same opponent indifferent interval (A, B) , by Lemma 35 we have $K(x, y) = K(x, x)$, and thus the K -function is P constant.

(\longleftarrow) This direction is immediate from Lemma 24. ◀

4.2 Characterization of Strong P Opponent Indifference

In this section, we will show how the strong version of P opponent indifference characterizes Sonas-like curves. We begin with an auxiliary lemma:

► **Lemma 36.** *If a rating system's K -function is constant over (A, B) and the skill curve is separable over (A, B) with a linear bisector, then the rating system is strongly opponent indifferent over (A, B) .*

Proof. Plugging the equations into the expected gain function for $x, x^*, y, y + \delta \in (A, B)$, we have

$$\begin{aligned} \gamma(x, x^* \mid y, y + \delta) &= K(x, y)(\sigma(x^*, y + \delta) - \sigma(x, y)) && \text{Lemma 15} \\ &= C \cdot (\beta(x^*) - \beta(y + \delta) + 0.5 - (\beta(x) - \beta(y) + 0.5)) \\ &= C \cdot (\beta(x^*) - \beta(x) - \beta(\delta)) \end{aligned}$$

which depends only on x, x^* , and δ . ◀

21:18 Opponent Indifference in Rating Systems

► **Definition 37** (Strong P Opponent Indifference). For $P \in (0, 0.5]$, a rating system is **strongly P opponent indifferent** if the rating system is strongly opponent indifferent over (A, B) for all A and B where¹¹

$$\sigma(A, B) > 0.5 - P.$$

► **Theorem 38** (Characterization of Strongly P Opponent Indifferent Rating Systems). A nontrivial continuous rating system (σ, α) is strongly P opponent indifferent if and only if the skill curve is P separable with a linear bisector and the K -function is P constant.

Proof. (\rightarrow) Assume the rating system is strongly P opponent indifferent. Because strongly P opponent indifferent rating systems are also P opponent indifferent, we already have that the skill curve is P separable and the K -function is P constant by Theorem 32. It only remains to prove is that the skill curve has a linear bisector.

The skill curve has a continuous bisector β by Lemma 34. Recall from Theorem 29 that, for any strongly opponent indifferent interval (A, B) , σ is separable on this interval with a bisector that is linear over the interval. Moreover, the slope of β on this interval must be $\frac{\sigma(B, A) - 0.5}{B - A}$.

From Lemma 33 and the fact that the rating system is nontrivial, we know there exists some $0 < \epsilon$ and $a \in \mathbb{R}$ such that $(a - \epsilon, a + \epsilon)$ is a nontrivial strongly opponent indifferent interval. Since the bisector is linear and nontrivial over $(a - \epsilon, a + \epsilon)$, we know $0.5 - P < \sigma(a - \epsilon, a + \epsilon) < 0.5$, and thus the slope of β is positive over $(a - \epsilon, a + \epsilon)$. Because of this, there must be some value c where $\sigma(a, c) = \sigma(a - \epsilon, a + \epsilon)$. Thus $(a - \epsilon, a + \epsilon)$ and (a, c) are overlapping strongly opponent indifferent intervals so β must have the same slope over both intervals. In order for this to be the case, $c = a + 2\epsilon$.

We can continue this argument, repeatedly adding ϵ , to show that β has the same slope as the interval $(a - \epsilon, a + \epsilon)$ for all $r > a$. We can also make a similar argument subtracting ϵ each time to show that the bisector has the same slope as the interval $(a - \epsilon, a + \epsilon)$ for $r < a$. Thus, the bisector's slope is constant. Since any vertical translation of β is also a bisector, there exists a linear bisector.

(\leftarrow) This direction follows directly from Lemma 36. ◀

► **Definition 39** (P Translation Invariant). A rating system (σ, α) is **P translation invariant** if for every $x, y \in \mathbb{R}$ where $\sigma(x, y) \in (0.5 - P, 0.5 + P)$, we have $\sigma(x, y) = \sigma^*(x - y)$ for some $\sigma^* : \mathbb{R} \rightarrow [0, 1]$.

► **Lemma 40.** A skill curve is *Sonas-like* with parameters a, s if and only if it is $\sigma(s, 0) - 0.5$ separable with $\beta(x) = ax$. (*Sonas-like skill curves are defined in Definition 3.*)

Proof. (\rightarrow) This direction follows from Definition 3, as $\sigma(s + x, x)$ is constant for all x .

(\leftarrow) From the definition of P separable and linear bisectors, for all $x, x + \delta$ where $\sigma(x, x + \delta) \in (0.5 - P, 0.5 + P)$, we have

$$\sigma(x, x + \delta) = \beta(x) - \beta(x + \delta) + 0.5 = -\beta(\delta) + 0.5,$$

and so the rating system is P translation invariant. Thus, for all $|x - y| < s$ we have $\sigma(x, y) = ax - ay + 0.5$. Since σ is continuous, it additionally holds that $\sigma(x, y) = ax - ay + 0.5$ when $|x - y| = s$. ◀

¹¹ Similar to P opponent indifference, this definition is equivalent to the one given in the introduction, despite looking slightly different (see Lemma 44 in Appendix A for details).

► **Corollary 41.** *A rating system is strongly P opponent indifferent if and only if it is Sonas-like with a P constant K -function.*

Proof. This follows directly from Theorem 38 and Lemma 40. ◀

References

- 1 Universal rating system™. URL: <http://universalrating.com/>.
- 2 Sharad Agarwal and Jacob R Lorch. Matchmaking for online games and other latency-sensitive p2p systems. In *Proceedings of the ACM SIGCOMM 2009 conference on Data communication*, pages 315–326, 2009.
- 3 Josh Alman and Dylan McKay. Theoretical foundations of team matchmaking. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, pages 1073–1081, 2017.
- 4 Todd Bardwick. Cscs informant – observations about chess rating distribution and progression. URL: <https://coloradomasterchess.com/informant-ratings-and-expectations/>.
- 5 Shuo Chen and Thorsten Joachims. Modeling intransitivity in matchup and comparison data. In *Proceedings of the ninth acm international conference on web search and data mining*, pages 227–236, 2016.
- 6 Zhengxing Chen, Yizhou Sun, Magy Seif El-Nasr, and Truong-Huy D Nguyen. Player skill decomposition in multiplayer online battle arenas. *arXiv preprint*, 2017. [arXiv:1702.06253](https://arxiv.org/abs/1702.06253).
- 7 Pierre Dangauthier, Ralf Herbrich, Tom Minka, and Thore Graepel. Trueskill through time: Revisiting the history of chess. *Advances in neural information processing systems*, 20, 2007.
- 8 Aram Ebtekar and Paul Liu. Elo-mmr: A rating system for massive multiplayer competitions. In *Proceedings of the Web Conference 2021*, pages 1772–1784, 2021.
- 9 Arpad E. Elo. The proposed uscf rating system, its development, theory, and applications. *Chess Life*, XXII(8):242–247, August 1967.
- 10 Vitalii Emelianov, Nicolas Gast, and Patrick Loiseau. Fairness in selection problems with strategic candidates. *arXiv preprint*, 2022. [arXiv:2205.12204](https://arxiv.org/abs/2205.12204).
- 11 RNDr Michal Forišek. Theoretical and practical aspects of programming contest ratings.(2009), 2009.
- 12 Mark E Glickman. A comprehensive guide to chess ratings. *American Chess Journal*, 3(1):59–102, 1995.
- 13 Kenneth Harkness. *The Official Chess Handbook*. McKay, 1967.
- 14 Ralf Herbrich, Tom Minka, and Thore Graepel. Trueskill™: a bayesian skill rating system. *Advances in neural information processing systems*, 19, 2006.
- 15 Stephanie Kovalchik. Extension of the elo rating system to margin of victory. *International Journal of Forecasting*, 36(4):1329–1341, 2020.
- 16 Lydia T Liu, Nikhil Garg, and Christian Borgs. Strategic ranking. In *International Conference on Artificial Intelligence and Statistics*, pages 2489–2518. PMLR, 2022.
- 17 Joshua E Menke, C Shane Reese, and Tony R Martinez. Hierarchical models for estimating individual ratings from group competitions. In *American Statistical Association*. Citeseer, 2007.
- 18 Tom Minka, Ryan Cleven, and Yordan Zaykov. Trueskill 2: An improved bayesian skill rating system. *Technical Report*, 2018.
- 19 SERGEY I Nikolenko and ALEXANDER V Sirotkin. Extensions of the trueskilltm rating system. In *proceedings of the 9th international conference on applications of fuzzy systems and soft computing*, pages 151–160. Citeseer, 2010.
- 20 Jeff Sonas. The sonas rating formula - better than elo? *ChessBase*, 2002. URL: <https://en.chessbase.com/post/the-sonas-rating-formula-better-than-elo>.
- 21 Jeff Sonas. The elo rating system – correcting the expectancy tables. *ChessBase*, 2011. URL: <https://en.chessbase.com/post/the-elo-rating-system-correcting-the-expectancy-tables>.

21:20 Opponent Indifference in Rating Systems

- 22 Jeff Sonas. What's wrong with the elo system? *ChessBase*, 2020. URL: <https://en.chessbase.com/post/what-s-wrong-with-the-elo-system>.
- 23 Weijie Su. You are the best reviewer of your own papers: An owner-assisted scoring mechanism. *Advances in Neural Information Processing Systems*, 34:27929–27939, 2021.
- 24 Weijie J Su. A truthful owner-assisted scoring mechanism. *arXiv preprint*, 2022. arXiv: 2206.08149.
- 25 u/vlfph. Farming volatility: How a major flaw in a well-known rating system takes over the gbl leaderboard. URL: https://www.reddit.com/r/TheSilphRoad/comments/hwff2d/farming_volatility_how_a_major_flaw_in_a/.
- 26 Lin Yang, Stanko Dimitrov, and Benny Mantin. Forecasting sales of new virtual goods with the elo rating system. *Journal of Revenue and Pricing Management*, 13(6):457–469, 2014.

A Extra Proofs

► **Lemma 42.** *If a function $f : (-A, A) \rightarrow \mathbb{R}$ is continuous and for $x, y \in (-A, A)$ we have $f(x + y) = f(x) + f(y)$, then f is linear over $(-A, A)$.*

Proof. To prove linearity, we need only prove $f(rx) = r \cdot f(x)$ for all $x, rx \in (-A, A)$. To do this we will prove the following propositions in order:

1. $f(nx) = n \cdot f(x)$ for n, x such that $n \in \mathbb{Z}$ and $x, nx \in (-A, A)$
2. $f(qx) = q \cdot f(x)$ for q, x such that $q \in \mathbb{Q}$ and $x, qx \in (-A, A)$
3. $f(rx) = r \cdot f(x)$ for r, x such that $n \in \mathbb{R}$ and $x, rx \in (-A, A)$.

The first proposition follows directly from the fact that $f(x + y) = f(x) + f(y)$, and it can be used to prove the second proposition. Let $q = \frac{a}{b}$ such that $a, b \in \mathbb{Z}$.

$$\begin{aligned}
 b \cdot f(qx) &= b \cdot f\left(\frac{a}{b} \cdot x\right) \\
 &= ab \cdot f\left(\frac{1}{b} \cdot x\right) && f\left(\frac{1}{b} \cdot x\right) \text{ is defined} \\
 &= a \cdot f\left(b \cdot \frac{1}{b} \cdot x\right) \\
 &= a \cdot f(x).
 \end{aligned}$$

Dividing both sides of the equation by b gives

$$\begin{aligned}
 f(qx) &= \frac{a}{b} \cdot f(x) \\
 &= q \cdot f(x).
 \end{aligned}$$

To show proposition 3, we can use proposition 2 and continuity. Since \mathbb{Q} is dense in \mathbb{R} , we know there exists a sequence of rationals q_n that converges to any $r \in \mathbb{R}$. Thus we can say

$$\begin{aligned}
 f(rx) &= f\left(\lim_{n \rightarrow \infty} q_n x\right) \\
 &= \lim_{n \rightarrow \infty} f(q_n x) && \text{continuity of } f \\
 &= \lim_{n \rightarrow \infty} q_n \cdot f(x) && \text{proposition 2} \\
 &= r \cdot f(x).
 \end{aligned}$$

So f must be linear over $(-A, A)$. ◀

► **Lemma 43.** *Definition 7 and Definition 31 are equivalent.*

Proof. (\rightarrow) Definition 7 is immediately stronger than Definition 31 since the former requires P -closeness among fewer numbers.

(\leftarrow) If a rating system satisfies Definition 31, then we may apply Theorem 32. Thus $\sigma(x, y) = \beta(x) - \beta(y) + 0.5$ and $K(x, y) = C$ for some constant C for x, y P -close. Looking at x, y and x^*, y both P -close for x, x^*, y as defined in Definition 7, we evaluate

$$\begin{aligned} \gamma(x, x^* \mid y, y) &= K(x, y)(\sigma(x^*, y) - \sigma(x, y)) && \text{Lemma 15} \\ &= C \cdot (\beta(x^*) - \beta(y) + 0.5 - (\beta(x) - \beta(y) + 0.5)) && \text{Theorem 32} \\ &= C \cdot (\beta(x^*) - \beta(x)) \end{aligned}$$

which only depends on the first two parameters, therefore satisfying Definition 7. \blacktriangleleft

► **Lemma 44.** *Definition 11 and Definition 37 are equivalent.*

Proof. (\rightarrow) Definition 11 is immediately stronger than Definition 37, since the former requires P -closeness among fewer numbers.

(\leftarrow) If a rating system satisfies 37, then we may apply Theorem 38. Thus $\sigma(x, y) = \beta(x) - \beta(y) + 0.5$ where β is linear and $K(x, y) = C$ for some constant C for x, y P -close. Looking at x, y and $x^*, y + \delta$ both P -close for x, x^*, y, δ as defined in Definition 7, we evaluate

$$\begin{aligned} \gamma(x, x^* \mid y, y + \delta) &= K(x, y)(\sigma(x^*, y + \delta) - \sigma(x, y)) && \text{Lemma 15} \\ &= C \cdot (\beta(x^*) - \beta(y + \delta) + 0.5 - (\beta(x) - \beta(y) + 0.5)) && \text{Theorem 38} \\ &= C \cdot (\beta(x^*) - (\beta(y) + \beta(\delta)) + 0.5 - (\beta(x) - \beta(y) + 0.5)) && \beta \text{ linear} \\ &= C \cdot (\beta(x^*) - \beta(x) - \beta(\delta)) \end{aligned}$$

which only depends on the first two parameters and the difference of the latter two parameters, therefore satisfying Definition 7. \blacktriangleleft