When it Pays to Think about the Competition, and When it Doesn't: Exploring Overconfidence Bias in Dynamic Games

Jennifer Cutler and Richard Staelin

The Fuqua School of Business
Duke University
100 Fuqua Drive
Durham, NC 27708 USA
E-mail: jennifer.cutler@duke.edu

Abstract Despite robust behavioral research that shows a widespread bias towards overconfidence in competitive scenarios, e.g., underestimating the competitor's skill level, there is little research on the long term costs associated with this bias. We develop a theoretical framework that allows us to explore systematic long-term ramifications of opponent skill estimation bias across different competitive contexts relevant to managers. We capture these contexts with dynamic branching games that are parametrized by four features. We use Monte Carlo estimation methods to test how the expected game outcomes compare under different types biases. The results suggest that bias in evaluating an opponent's skill level is less harmful when the opponent is more skilled, and when there is greater first-mover advantage. Furthermore, they suggest that if there is any effort cost associated with making a decision, then a bias towards overestimating the opponent's skill is never advantageous, while a bias towards underestimating can be advantageous in many contexts.

Keywords: bounded rationality, overconfidence bias, heuristics.

1. Introduction

Many firms invest heavily in competitive analysis, and proverbial advice such as "never underestimate the competition" abounds in the popular business literature. Yet, little has been done to formally test this stylized wisdom from a cost-benefit perspective. There is a prevailing intuition that being caught off guard by a competitor who is more skilled than expected has negative consequences— but what are the relative consequences of the opposite mistake? How much potential surplus would be lost by routinely assuming the competition to be more skilled than it actually is? We build a framework for exploring bounded rationality in dynamic branching games and use it to compare the relative payoff ramifications of different types of biases across different competitive contexts relevant to managers.

2. Background

Traditional game theory models have generally assumed that all players are fully rational. Recent research streams have adapted models of decision making to fit the more realistic assumptions of bounded rationality, where players act rationally within the bounds of their constraints on information processing (see Narasimhan et. al., 2005).

Two primary characteristics found in bounded rationality models are that players have different skill levels (which can lead to uncertainty about the skill levels of their opponents) and that the effort of optimizing decisions is costly. Researchers have used paradigms such as level-k and cognitive hierarchy models to incorporate heterogeneous reasoning levels into game theoretic models for some time (e.g., Camerer et al., 2004; Stahl and Wilson, 1995). These models both help explain observed non-equilibrium behavior in games and estimate empirical distributions of reasoning abilities and beliefs about others' abilities. We take a different approach: knowing that error is possible (and likely unavoidable) when boundedly rational managers form beliefs about the skill levels of the opponents they encounter, we explore, from a theoretical standpoint, the payoff ramifications associated with such errors, including relative cost of effort. How much does a manager stand to lose by repeatedly overestimating or underestimating the skill levels of the different opponents he encounters? Which features of games make one type of error better than another?

Our approach is similar to that of researchers such who have compared average outcomes produced by different strategy heuristics for individual decision makers when time or effort is costly or limited (e.g., Gabaix and Laibson, 2000; Johnson and Payne, 1985; and Payne et al., 1996). These researchers employed Monte Carlo simulation techniques to approximate the mean outcomes of different heuristic strategies when applied over a large set of normal or extensive form payoff arrangements that were generated in part using random number generation to capture the inherent variation in real-world situations. In all these studies, it was concluded that simplifying heuristics can perform better than rote optimization when the decision makers are subject to some elements of bounded rationality. We build on similar conceptual and methodological foundations, but add the complexity of a second strategic player and bias surrounding comparative skill.

3. Model Framework

3.1. Game Structure

Our first goal is to define a highly generalizable game structure that captures the critical features of strategic competitive interactions faced routinely by boundedly rational managers in a variety of settings without overfitting to a particular circumstance. As such, we look for general game characteristics rather than a unique example.

We observe that when navigating complex long-term competitive relationships, managers are often faced with decisions that have not only immediate payoff ramifications, but also affect the choices (and associated payoffs) that will be available in the future. For example, a firm may be deciding between releasing a newly developed product into the market at a low price to encourage trial, or at high price in order to obtain surplus from enthusiastic early adopters (e.g., Apple releasing the initial Ipod). There are immediate profits to be gained from the decision (in terms of profits generated from initial sales) but these different moves may also have longer-term strategic consequences, based in part on the actions taken by the firm's major competitor. For example, this competitor might be developing a competing product (e.g., Microsoft releasing the Zune), and after observing how the first firm priced its product, the competitor will decide whether to release its product at a regular or a substantially discounted price. When this happens, there will be an immediate shift in profits for both firms, and also set up the first firm for a counter-response in terms of adjusting prices or promotions or investing in new development efforts. If the first firm is myopic, then it will release its product at whatever price initially maximizes sales, without considering what the competition will do in response. If the firm is more sophisticated, it will consider possible competitive responses, its own counter-response, and so on, and make its initial decision with those downstream implications in mind. Of course, there is a limit on how far out even the most sophisticated manager can think, especially when the competitive horizon is long or complex.

We use a branching tree structure to represent such real world strategic decision situations where managers repeatedly make moves and counter moves. Thus, the players alternate making decisions over time, and each node in the tree has a payoff associated with it (which represents the value of the state of affairs at that moment in time, e.g., current market share or profits). To capture different skill levels, we let players vary in the time horizon over which they can optimize from any given moment in the game.

This leads us to assume two-player alternating-move finite-horizon games of perfect and complete (within one's foresight horizon[1]) information, where for parsimony we limit our attention to constant-sum payoffs. This structure has appeal for many reasons. Dynamic games allow foresight horizon to be used as a precise measure of skill, capturing varying degrees of myopia that managers employ within their long-term competitive landscapes. Games in which information is perfect and complete (within a player's foresight window) provide an excellent platform for studying the research question at hand, because variation in outcome is caused solely by variation in skill and beliefs. Finite horizon games apply to situations in which managers are competing over a set term (for example, sales over a holiday season, performance bonuses over a fiscal year) or when players make several moves that lead to a long-term stabilization of market shares. Many competitive marketing situations are inherently constant sum (for example, employees competing over a fixed bonus pool, firms competing for market share in an inelastic market). Even if the games are not constant sum they can be re-framed as such if the payoffs are normalized to reflect relative competitive advantage. For convenience, we restrict all payoffs to the range of 0 to 1, which corresponds nicely (but not restrictively) to market share. By convention, P1 always seeks to maximize payoffs, and P2 seeks to minimize payoffs (i.e., P2 seeks to maximize one minus the payoff).

We note that our game structure assumes state values (interim payoffs) at each non-terminal node. Though ubiquitous in computer science (e.g., see Hsu, 2002 and Russell and Norvig, 2003), which as a field is inherently concerned with bounded rationality, this notion of pre-terminal state values is not widely used in the economics and management game literature. Instead, the games studied previously are generally defined in term of the final (terminal) payoffs, a paradigm appropriate for managers with unconstrained reasoning abilities. We break from this literature since we believe interim payoffs capture real-world phenomenon applicable to boundedly rational managers. Intuitively, if a manager's forsight constraint precludes him from anticipating payoffs all the way out to the end of the game, he must use something observable in the interim to guide his behavior. For example, employees competing for a year-end bonus allotment may use quarterly performance reviews as a means of

[1] i.e., players can see all payoffs that fall within their foresight horizons perfectly

evaluating standing along the way. Or, politicians campaigning for an elected office may use opinion polls as a measure of vote market share at various intervals before the election takes place. Thus, the interim payoff values we preserve in our game structure can be considered to be imperfect assessments of the strategic advantage of a state in the game (i.e., of the final outcome to which the state will lead). Of course, there are numerous ways to evaluate the worth of such interim states, and devising a sophisticated method of doing so is a skill in itself. Our correlation parameter, ρ, which will be discussed in the next section, controls the reliability of these signals in our model.

Many well-known strategy games fit this structure, including Go and Chess, which are both renowned for their strategic complexity and played in highly competitive international tournaments. Like all games of complete and perfect information, the process of solving Go or Chess is theoretically trivial—but the branching game trees are so complex that not even the best computers in the world can model the full game tree. Artificially intelligent players instead must be forward-looking to build a game tree that extends as many rounds into the future as they are capable of, and then optimize play over that limited horizon. Final scoring rules or other evaluation functions are applied to transient states of the board as a way of assessing the value of the board, even though only the configuration present when the game ends is used to tabulate final score (see Russell and Norvig, 2003).

By following in this tradition, we are able to explore player navigation in a class of games that has real-world relevence for manager but for which the curse of dimensionality precludes calculation of classically rational behavior in all but the simplest of cases. Strategic decision making under constrained foresight can be modeled in even the most complex of games, as it removes the link between game complexity and calculation complexity.

3.2. Game Parametrization

As we are interested in looking at a variety of different games that might be encountered by players in the real world, we abstract away the labels of the actions that define the circumstances of any particular game, reducing each game to a pattern of branches and payoffs (this is similar in method to the generalized decision structures used by Gabaix and Laibson, 2000; Johnson and Payne, 1985; and Payne et al., 1996) that can be populated by any number of possible payoff arrangements. Of course, real world games vary in the patterns they create: one game will have one arrangement of payoffs (based on its set of actions) and another game will have quite a different set of payoffs. Thus we create *sets* of games that include all possible games a player could encounter that fit certain useful pattern definitions. By showing the effects of bias conditions over entire game sets, we attempt to model the overall effects of such biased beliefs applied consistently over many situations with many different opponents. We allow the number of periods in the game tree to vary by length (L). For simplicity, we restrict L to even integers, so that both players have an equal number of moves. In this initial analysis, all decision nodes contain a choice between exactly two actions. However, the model can be extended in future work to allow such complexity descriptions to vary.

The payoff correlation parameter ρ is used to define the relationship between early interim payoff signals and their downstream consequences. We begin by allowing the initial period interim payoffs to be drawn randomly from some defined distribution. In our case, we use the uniform distribution $U(0, 1)$, but this can easily be modified within our framework. This initial distribution is used to capture the inherent variation of possible circumstances. The values of subsequent nodes are determined such that $v_t = \rho v_{t-1} + (1-\rho)\epsilon_t$ where v_t is the value associated with a node at period t, v_{t-1} is the value of the preceding node on the same path (parent node), ϵ_t is a draw from $U(0,1)$, and $0 \leq \rho \leq 1$. In this way, ρ captures the strength of the signal interim payoffs provide for future payoffs.

3.3. Player Parametrization

Behavioral research has shown that managers engage in a process of limited-horizon reasoning when engaging in dynamic strategic behavior (e.g., Camerer and Johnson, 2004; Johnson et al., 2002; Stahl and Wilson, 1995). Rather than performing full backward induction, players look forward a certain number of periods, and optimize only over that window, as if the game truncated there. Players vary in the number of future periods they can think though. These empirical finding drive the inclusion of such behavior in our model.

In line with extant level of thinking models, we use the parameter k_i to represent player i's skill, and specifically define k_i as *foresight horizon*: the number of periods out into the future he can think through, including the present period. Players also have beliefs about their opponent's skill level: we define parameter b_{ij} as player i's belief about k_j. A player's skill k_i and belief b_{ij} map to a player type that chooses his action at each period by considering only the subgame that starts at the current period and ends k_i periods later (even if the true game length extends beyond that horizon). In other words, at each period in which he makes a decision, player i (P_i) creates a truncated version of the game tree that starts at the current node (period t) and ends after k_i periods. From P_i's point of view, he cannot think beyond period $t+k_i-1$ and thus this player's objective is to optimize his outcome in period $t+k_i-1$ (effectively acting as if the game ends at that point, or, equivalently, that all future states that exist beyond this period have the same payoff as the their preceding state)[2]. Our motivation in this type definition is to capture the heuristic-like behavior of players under cognitive constraints.

To describe the decision rule in more detail, we break beliefs into two categories: $b_{ij} \geq k_i-1$ (i.e., P_i believes his opponent can see out at least as far in the game as he himself can at any particular period[3]) and $b_{ij} < k_i - 1$ (P_i believes his opponent is more than one level less skilled than he himself is). When $b_{ij} \geq k_i-1$, P_i determines his move based on full backward induction over this truncated game. This is because at each period within this truncated game, P_i believes that the player in control can see through the entire truncated game. Note that all beliefs b_{ij} that are greater than or equal to $k_i - 1$ map to the same decision algorithm for P_i. This is because a player constrained by his own foresight horizon cannot anticipate what his opponent would do outside that window, even if he believes his opponent to be looking farther (and behavioral research supports that people do behave as if the game ends at the

[2] P_i optimizes for period $t+k_i-1$, rather than period $t+k_i$, as we have defined k_i to include the current period.

[3] The reason $k_i - 1$ rather than k_i is the critical value is due to the sequential nature of the game. If P_i is making a decision in period 1, and is a level 4 thinker, then he can "see through" to period 4, and nothing beyond that. Since his opponent P_j will not be making a move until the next period, P_j can be allotted a maximum of 3 levels from P_i's perspective (thus reaching period 4) before $P_i's$ own constraint prevents him from seeing father, even through P_j's eyes.

limit of their foresight horizon, even if they are aware that it doesn't and that other players might be thinking farther, e.g., Johnson et al., 2002). Thus, as long at as the more limited skill player knows his opponent has a greater skill level than he does (or, more precisely, not more than one level less), fine-tuning his estimate of his opponent's skill will not change his behavior. However, the higher skill level player needs to make a more precise estimate of the less skilled player's skill level.

On the other hand, when $b_{ij} < k_i - 1$, P_i must first consider the even smaller subgames he believes his opponent (P_j) will be optimizing over at each of the periods P_j controls that fall within P_i's horizon. P_i must do this first, since he believes P_j will be using different critera than the final period values of P_i's truncated game. Only then can P_i backward induct over his full foresight horizon to determine his own optimal move, given his beliefs about P_j's less sophisticated decisions. More specifically, for each state that could occur in period $t+1$ (as a result of P_i's choice in period t), P_i creates a subgame that commences with that state and extends to a length of b_{ij}. He backward inducts over each of these subgames to determine what P_j will choose if found in that state in the second period. Then for each of the potential states that could exist in period $t+1$, P_i can prune off the rejected options, as well as all the downstream branches of the tree that stem from these rejected second period options. P_i then repeats this process for each period that P_j controls within his foresight window, further pruning the tree each time. When he has determined what P_j will decide at all decision sets in the foresight window, he can backward induct over the pruned tree, and arrive at his optimal choice for the current period.

For simplicity, we define P1 to always be the more skilled player, such that $k_1 > k_2$. Furthermore, we do not consider cases in which P_i's beliefs about P_j's beliefs about P_i (i.e. P_i's belief about b_{ji}) vary enough to change the decision algorithm described above. For example, if P_i has belief $b_{ij} = 3$, then the optimal backward induction process P_i uses over the subgames of length b_{ij} (to determine what P_j will choose at a given decision set) would change if P_i believed that P_j believed that P_i was a level 1 thinker vs. a level 2 thinker (note, however, that as per the above reasoning, all beliefs of this nature that are greater than or equal to 2 map to the same full backward induction over the subgame—so it is only when $b_{ij} \geq 3$ that these higher order beliefs could affect decision processes). For simplicity, we do not consider variation in higher order beliefs, and assume that, for any opponent skill belief b_{ij}, P_i performs a full backward induction when hypothetically selecting moves from P_j's perspective.[4] We do this because there is limited room for such variation within shorter foresight horizons and experimental work suggests that people generally have quite limited levels of thinking (see, for example, Camerer et al., 2004 and Stahl and Wilson, 1995).

For simplicity, we assume that each player holds a single value for b_{ij} and applies it with certainty (i.e., instead of specifying a distribution of possible opponent beliefs, or updating beliefs over time). This assumption is in line with behavioral work that finds individuals show overconfidence about their judgments (e.g., Hoffrage, 2004), and is further justified by the limited opportunity for learning within the scope of the model. Players are not engaging in long, repeated games with familiar opponents, but rather engaging in a large variety of short, novel one-

[4] However, preliminary findings suggest that relaxing this assumption only strengthens the results found in this paper.

time games with novel opponents and only stochastic feedback. Empirical findings show that people are slow to update existing beliefs, especially amid noisy signals, and that when they do, they overweight prior beliefs (e.g., Boulding et al., 1999; Camerer and Lovallo, 1999). Furthermore, the first move, which must occur before any learning is possible, has the most influence on the game, especially when $\rho > 0$. Thus we believe our assumed model captures the general belief conditions found in many sequential games.

We also note that we use foresight horizon as our singular measure of player skill. There are, of course, many dimensions to skill when playing games in real life. For example, players may differ on their ability to accurately assess payoffs or to correctly apply backward induction. For parsimony in our manipulation, we create a level playing field on all dimensions except foresight horizon. This follows in the tradition of many analytic and level-k models, and allows a clean manipulation of skill levels. However, the model can be extended in future research to account for other such dimensions of skill—for example, by allowing players to “see" only a variably imperfect correlate of the true payoff for any given state.

Finally, we add for clarification that we assume that players do not alter their decision rules based on the value of ρ. This is because we are attempting to capture the overall effect of different bias conditions applied heuristically by managers repeatedly over many different contexts and against many different opponents. The value of ρ in a real-world setting would be difficult to observe precisely, especially under constrained reasoning. Furthermore, knowledge of ρ would only potentially change players' decision rules if players were endogenously concerned with minimizing effort cost. We assume that this is not the case and will explain our reasons for this in Section 3.5..

3.4. Bias Conditions

We begin our theoretical experiment by creating three bias conditions (B) that correspond to three “worlds" in which players all exhibit one of three types of bias in estimating their opponent's skill level, and are unaware that the bias is present.[5] We first consider the accurate ($B = A$) opponent skill estimation condition, where $b_{ij} = k_j$, to model a condition of no bias. This scenario is most similar to traditional game theoretic methods. Players have exogenously defined skill levels (where $k_1 > k_2$), and we assume that these skill levels are common knowledge.

We next consider an opponent underestimation ($B = U$) condition, corresponding to overconfidence bias. To model a slight population-wide opponent underestimation bias, we set $b_{ij} = k_j - 1$, such that each player believes his opponent to be exactly one level less skilled than is actually the case. Guided by the observation that people are often unaware of bias, we initially assume that both players think their beliefs map to true values that are common knowledge. P1 thinks that k_1 and $k_2 - 1$ are the common knowledge skill levels for P1 and P2, while P2 thinks that $k_1 - 1$ and k_2 are the levels for P1 and P2 that are common knowledge. However, we soon show in Lemma 2 that this assumption can be relaxed.

Finally, we test a condition in which players exhibit a bias towards overestimating opponent skill ($B = O$). Though such “underconfidence" does not appear to be prevalent in nature, this test provides a useful theoretical tool for understanding

[5] We later show that neither awareness of the bias nor asymmetric bias will affect player decisions or game outcome.

overconfidence. This condition is defined similarly to the U condition, except that now both players believe their opponent to be exactly one level *more* skilled than is actually the case ($b_{ij} = k_j + 1$). P1 thinks that k_1 and $k_2 + 1$ are the skill levels for P1 and P2 that are common knowledge, while P2 thinks that $k_1 + 1$ and k_2 are the skill levels for P1 and P2 that are common knowledge.

A logical implication of this framework that simplifies our analyses is that, since P2's actions are the same for all beliefs $b_{21} \geq k_2 - 1$, and since $k_1 > k_2$, then the two bias conditions in which players misestimate opponents skill level (O and U) have no effect on P2's decision process compared to the accurate estimation condition (A). Even if P1 is only one level more skilled than P2, an underestimation bias on the part of P2 will yield $b_{21} = k_1 - 1 = k_2$, which maps to the same decision in all cases as accurate estimation ($b_{21} = k_1 = k_2+1$) and overestimation ($b_{21} = k_1+1 = k_2+2$). It is only when the *more skilled* player has biased beliefs that the game outcome is potentially changed. This is because only the more skilled player has the opportunity (capacity) to think further out than his opponent and this foresight can result in changes in his decision process. This leads to the following lemma.

Lemma 1. *The weaker player's bias condition has no effect on game outcome*

Thus, throughout our analysis, we focus on the implications of opponent estimation error from P1's perspective. All results will hold whether or not the weaker player exhibits bias.

Lemma 1 implies that P2's bias doesn't affect his own behavior. Consequently, P1's *awareness* of P2's bias does not change P1's behavior, since P2 will behave the same in all conditions. Nor will P2's awareness of P1's bias change P2's behavior. To see this, note that when P2 is looking hypothetically through P1's eyes over his truncated subgames, trying to anticipate what P1 will do, he only has $k_2 - 1$ periods to work with—which leaves only $k_2 - 2$ periods in which he can anticipate what P1 thinks about P2. In other words, all of P2's beliefs about what P1 believes about P2 that are $\geq k_2 - 2$ map to the same decision process and the same outcome. In bias conditions U, A, and O, respectively, P1 believes k_2 to be $k_2 - 1$, k_2, and $k_2 + 1$. All of these values are $\geq k_2 - 2$, and thus even if P2 is aware of P1's bias, it does not change his behavior, due to the limitations of his own cognitive constraints.

Lemma 2. *The outcome of the game is the same whether or not players are aware of the other player's bias.*

Note that Lemmas 1 and 2 hold in our model because we are looking as small errors in opponent estimation, or slight biases. If we explored conditions in which players underestimate each others' skill to large degrees (specifically, with an underestimation bias that is $> 1 + k_1 - k_2$) then player decisions would potentially be affected. Together, these lemmas show that our model holds under assumptions that are less restrictive than we initially laid out. For logical equivalence, we need only assume that P1 is more skilled than P2 (i.e., $k_1 > k_2$)[6], that P2 believes P1 can see out at least as far as he can (i.e., $b_{21} \geq k_2 - 1$), that P1 will not underestimate P2 by more than three levels[7], and that both players are aware of all these

[6] To allow "room" for P1 to effectively overestimate P2, we focus on cases in which $k_1 > k_2 + 1$.

[7] This is the minimum lower bound over all conditions for which P2's beliefs about P1's beliefs about P2 might change P2's decision behavior.

things. In addition, to test out bias conditions, we further assume that P1 exhibits the following condition-dependent beliefs. For $B = U$, P1 underestimates P2's skill level such that $b_{12} = k_2 - 1$. For $B = A$, P1 accurately estimates P2's skill level such that $b_{12} = k_2$. And for $B = O$, P1 overestimates P2's skill level such that $b_{12} = k_2 + 1$.

3.5. Effort Cost

Under boundedly rational paradigms, it is generally accepted that there is some cost of information acquisition and processing, whether it be opportunity cost, error introduction, or sheer disutility of effort (e.g., see Shugan, 1980). With this noted, there is no consensus among researchers on how to precisely define such a cost function for decision making. Moreover, in real life these costs are likely to be highly variable across persons and contexts. Consequently, to maintain greatest external validity, we refrain from defining any specific cost functions, and assume only that such effort cost strictly increases with the amount of information processed, which can be represented through the number of game tree nodes a player generates when making a decision. This is consistent with behavioral traditions that use elementary information processes as a measure of the cost of cognitive effort (see Newell and Simon, 1972 and Payne et al., 1995). We use $\mathbb{C}(B)$ as the cost of effort exerted by P1 under bias condition B for a set k_1 and k_2 where the cost is some increasing function of nodes examined during the first move.

We also assume that players do not account for any effort cost when making decisions. Rather, players' decision rules are driven by their skill levels and beliefs, and we compute comparative effort costs post-hoc, to show the relative long-term advantages of different heuristics applied by decision makers automatically (e.g., see Stahl, 1993). This assumption is perhaps unusual in the game theory literature, especially when costs are considered purely search costs, rather than optimization costs. However, we follow in the tradition of bounded rationality paradigms and take our players' decision rules to be heuristic-like (based on empirical observation) rather than strict optimization. Thus, the players within our model apply the same decision rules (based on their own skill and their beliefs about their opponent's skill) to each decision they encounter, as a matter of course, without considering effort cost. This approach is in line with empirical and theoretical work on heuristics and biases and circumvents a general problem with optimization in bounded rationality paradigms when decision effort itself costly: the act of cost-aware decision optimization becomes impossible, as the decision of how to decide how to decide (and so forth) becomes an infinite regress where each such higher order decision exacts its own cost (e.g., see Gigerenzer and Selten, 2001).

3.6. Net Expected Outcome

We define *net expected outcome* as the gross expected outcome for a given bias condition and parameter set minus the cost of implementing the decision rule. We use $\mathbb{E}_{\mathbb{N}}(B)$ to refer to the net expected outcome for bias B over a fixed game and skill set, where $\mathbb{E}_{\mathbb{N}}(B) = \mathbb{E}(B) - \mathbb{C}(B)$. Of course, we cannot compute meaningful numeric values for $\mathbb{E}_{\mathbb{N}}(B)$ without assigning a a shape and relative scale to the cost function, which is outside the scope and motivation of this paper. However, with the one assumption that cost is a strictly increasing function of information searched, we can determine useful ordinal properties and trends. Thus, in our analysis we will look for insights relating to the the relative net expected outcomes of the three bias conditions as a function of the game parameters L and ρ and the player parameters k_1 and k_2. Figure 1 illustrates the structure of the model framework.

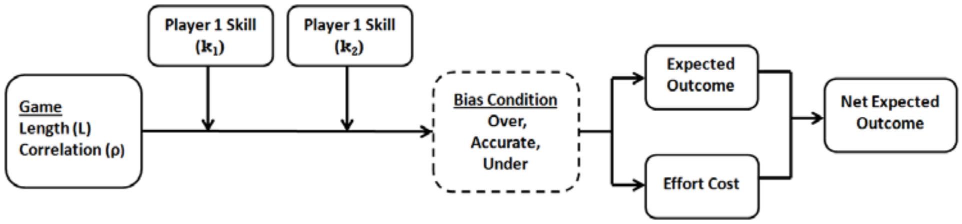

Fig. 1: Model Framework

4. Analysis

4.1. Estimating Expected Outcomes and Costs

For any specific game tree, the unique outcome of the game is entirely predictable given k_1, k_2, and bias condition B. Unlike full rationality models, the players themselves may not accurately predict final outcomes while in early periods, and may make mistakes (when foresight is incomplete or there is estimation error) that lead the game into a state they did not anticipate. However, the outcome that will be arrived at is deterministic from an outside perspective, i.e., by someone who can see the entire game tree and knows k_1, k_2, and B, and can apply the appropriate decision algorithms from the perspective of each player to arrive at the final outcome. As the outcome for any single game is of little generalizable value, we focus on estimating the expected values for the outcomes over infinite sets of possible game trees that make up game sets defined by specific values of L and ρ with player skills defined by k_1 and k_2. We use the notation $\mathbb{E}(B)$ as a shorthand for the expected outcome under bias condition $B \in \{O, A, U\}$ for a fixed parameter set.

We use Monte Carlo simulation to estimate the expected outcomes of each bias condition over a range of parameter values that we believe are reasonable within our general premise of bounded rationality. We test all parameter combinations that meet the requirements of our model that fall within the range $2 \leq k_2 \leq k_1 - 2 \leq k_1 \leq L \leq 10$ and ρ in 0.1 increments from 0 to 1 (i.e., $\rho = \{0, .1, .2, .3, ..., 1\}$).

We use the Visual Basic environment in Microsoft Excel to write a program that builds sample game trees of any specified L and ρ, where the payoffs in the first period (v_1) are drawn randomly from $U(0, 1)$ using Excel's random number generator, and each child node's payoff is equal to its parent node's value plus error weighted by ρ such that $v_t = \rho v_{t-1} + (1 - \rho)\epsilon_t$, where ϵ_t is a draw from $U(0, 1)$. The decision rules described in Section 3.3. are programmatically applied to find the unique outcome for each generated game tree according to any specified k_1, k_2, and B.

We generate two thousand game trees for each unique combination of game parameters (L and ρ) included in the range defined above. For each tree generated, we record the outcome reached under each bias condition (B) for each combination of player skill parameters (k_1 and k_2) included in the range defined above. The mean outcome over all trials of a unique parameter set is used to approximate the expected outcome for that parameter set. For simplicity, we average the results for consecutive values of k_2 to eliminate main effects of the parity of k_2 which we do no expect to be useful in terms of managerial insights. We also count the number of nodes searched in the first turn of P1 to use as a basis for ordinal comparisons of effort cost (which as defined in Section 3.5., is considered to be some (any) strictly increasing function of information searched.

4.2. Results

Comparing the estimates for all combinations of parameter values in the range tested, we find the following results.

Result 1. *The expected outcome is greatest when P1 accurately estimates his opponent's skill level, followed by when he overestimates, followed by when he underestimates, for all otherwise fixed parameter values;* $\mathbb{E}(A) > \mathbb{E}(O) > \mathbb{E}(U)$ *for all fixed* k_1, k_2, L, *and* ρ.

Note that if the parameter range allowed $k_2 \geq k_1$ then the inequalities would become weak, as the search behavior (and thus outcomes) would be the same in all three conditions from P1's perspective.

Result 1 supports conventional wisdom and intuition. However, we expect that the magnitude of disparity between the raw exceptions changes in different settings. Thus, we next investigate the relationship between opponent skill and the differences in $\mathbb{E}(B)$. We first note that an increase in k_2 when $k_1 - k_2$ is held constant implies that both players are getting more skilled, while the disparity between them stays the same. On the other hand, an increase in k_2 while k_1 is held constant implies that P2 is getting more skilled, while the disparity between the players diminishes. In both cases, we find that increasing k_2 decreases the differences in $\mathbb{E}(B)$.

Result 2. *The expected outcomes of all bias conditions converge as both players become more skilled together; differences in* $\mathbb{E}(B)$ *decrease with* k_2 *for any fixed* $k_1 - k_2$, L, *and* ρ.

Result 3. *The expected outcomes of all bias conditions converge as P2 becomes more skilled relative to P1; differences in* $\mathbb{E}(B)$ *decrease with* k_2 *for any fixed* k_1, L, *and* ρ.

To explain the intuition for this result, we note that as P2 becomes more skilled, P1 has less control over the outcome, regardless of bias condition. The longer the horizon over which P2 has full foresight, the greater P2's ability to influence the outcome, which limits the range of possible outcomes available to P1. Thus the range between the worst outcome and best outcome for P1 decreases when he is playing against more skilled opponents. As a result, there is less relative payoff decrease to P1 as a result of error. This is true both as P2 becomes more skilled in absolute terms, as well as relative to P1. As P1's skill advantage over P2 decreases (i.e., as the players become more evenly matched, regardless of absolute skill level), there are fewer periods over which P1 can use his advantage—and thus fewer periods in which estimation error can detract from the potential outcome. As a result, the differences in $\mathbb{E}(B)$ decrease whenever average opponent skill increases, regardless of whether both players are getting more skilled together, or whether the disparity between them is decreasing.

As expected, we also find that the expected outcomes of each of the three bias conditions converge as ρ increases.

Result 4. *The difference between the expected outcomes of overestimation and underestimation decrease with payoff correlation for any fixed game and skill set.*

This is intuitive because increasing ρ increases the advantage generated by the first move, thus decreasing the influence P2's strategy will have on the outcome, which thereby decreases the expected payoff loss associated with opponent estimation errors.

In addition to having asymmetric effects on expected payoffs, we also fine that different types of errors in estimating opponent skill also have asymmetric implications for effort cost. When underestimating his opponent, P1 can prune off much of the game tree without ever having to generate of process the payoffs associated with those states. This results in an effort cost savings. This is true to a less extent with accurate estimation, and to an even less extent with over estimation.

Result 5. *When a player's own level of thinking is fixed, it costs him the most to overestimate his opponent, less to accurately estimate, and least to underestimate;* $\mathbb{C}(O) > \mathbb{C}(A) > \mathbb{C}(U)$ *for any fixed* k_1 *and* k_2.

Note that the strict inequality holds for the parameter range tested. If the the range allowed $k_2 \geq k_1$, then the inequalities would become weak, as the search behavior (and thus costs) would be the same in all three conditions from P1's perspective.

From Results 1 and 5 we know that $\mathbb{E}(O) < \mathbb{E}(A)$ and $\mathbb{C}(O) > \mathbb{C}(A)$ for all game and player sets that fit the requirements of our model. From here one can directly conclude that $\mathbb{E}(O) - \mathbb{C}(O) < \mathbb{E}(A) - \mathbb{C}(A)$. In other words, overestimation always has a lower expected net (as well as gross) return than accurate estimation, regardless of game or player parameters. In fact, the difference between the net expected outcomes is necessarily larger than the difference between gross expected outcomes. This leads us to Result 6.

Result 6. *The expected net payoff for overestimation is always strictly less than the expected net payoff of accurate estimation for any fixed parameter set:* $\mathbb{E}_{\mathbb{N}}(O) < \mathbb{E}_{\mathbb{N}}(A)$.

There is no such strict dominance with underestimation. From Results 1 and 5, we know $\mathbb{E}(U) < \mathbb{E}(A)$ and $\mathbb{C}(U) < \mathbb{C}(A)$, from which we cannot determine a general ordinality for $\mathbb{E}(U) - \mathbb{C}(U)$ versus $\mathbb{E}(A) - \mathbb{C}(A)$. The difference in expected net returns of under vs. accurate estimations will depend on the magnitudes of each term, which are determined by the specific parameters of a game and player set, as well as the specific cost functions used. It is not our goal in this paper to propose valid cost functions. Still it is possible to gain additional insights by considering the relationship of the parameters ρ and k_2 on the differences in $\mathbb{E}_{\mathbb{N}}(B)$.

Considering first the minimum values for k_2 and ρ, we note that, depending on the cost function, $\mathbb{E}_{\mathbb{N}}(U)$ can take any of three positions: it can be less than $\mathbb{E}_{\mathbb{N}}(O)$, it can be greater than $\mathbb{E}_{\mathbb{N}}(O)$ but less than $\mathbb{E}_{\mathbb{N}}(A)$, or it can be greater than $\mathbb{E}_{\mathbb{N}}(A)$. However, as either k_2 or ρ increases, the expectation disadvantage of U decreases relative to A and O. Regardless of where $\mathbb{E}_{\mathbb{N}}(U)$ begins relative to $\mathbb{E}_{\mathbb{N}}(O)$ and $\mathbb{E}_{\mathbb{N}}(A)$, the slope differences will cause there to be some critical value of both k_2 and ρ, above which $\mathbb{E}_{\mathbb{N}}(U)$ is the best performing condition if the trend lines are extrapolated. As we are not defining cost function scales, we cannot say if the critical value will occur within the parameter limits imposed by a player's own skill constraint. However, we can say that this becomes more likely that the cost of information acquisition and processing increases.

This brings us to Result 7.

Result 7. *Underestimation can yield the greatest net expected outcome for a parameter set. This is more likely to occur when the opponent is highly skilled, when the first mover advantage is strong, and/or when effort costs are high;* $\mathbb{E}_{\mathbb{N}}(U)$ *becomes more likely to to be higher than both* $\mathbb{E}_{\mathbb{N}}(O)$ *and* $\mathbb{E}_{\mathbb{N}}(A)$ *as* k_2, ρ, *and effort costs increase.*

5. Discussion

The ultimate goal of this paper is to provide an initial exploration into the question of if and when overconfidence can be beneficial to managers who make frequent complex competitive business decisions. In order to do this we needed to develop a new and general framework for analyzing boundedly rational players in "large world" (Savage, 1954) complex games. Our framework uses a branching decision tree with interim payoffs to represent a strategic game between two players where players make sequential moves over time and have limited foresight. Because our model includes a skill constraint for each player, we are able to explore branching game structures that have real-world applicability but that can be quite difficult to manage under traditional assumptions of rationality due to the curse of dimensionality. Given our interest in generalizable conclusions, we build a Monte Carlo simulation program to estimate the expected payoffs over a the distribution of possible payoff structures associated with any given game length, payoff correlation, players' skill levels, and players' beliefs about their opponent's skill level. We believe this framework could be useful to others interested in bounded rationality and branching sequential games.

Our results suggest that bias in evaluating an opponent's skill is less harmful to expected payoff when the opponent is more skilled, and when there is greater first-mover advantage. Furthermore, they suggest that if there is any effort cost associated with the making a decision, then a bias towards overestimating the opponent's skill is never advantageous, while a bias towards underestimating can be advantageous in many contexts. Thus, the overconfidence bias behavioral researchers have observed in the population may actually be helpful, rather than detrimental, as is often suggested, and we provide initial insight into when this is more likely to be the case.

Although these initial theoretical experiments begin to shed light on the relationships of interest, they have several limitations, and thus should be considered as only a start to understanding the greater relationships between skill constraints, opponent estimation errors, and outcomes. For example, the game contexts the model framework covers in this initial exploration are limited, but could, in the future, be extended to capture non-zero-sum games, games with greater complexity, and state-dependency in the parameters. In terms of players, the extant model considers only one dimension of skill (foresight horizon) and thus the results do not generalize to estimating other dimensions of competitive capabilities, such as sophistication in estimating the interim payoffs (or resources of the competing firm). We believe future research can build off this framework to address these and other limitations, thereby providing deeper and broader insights to advise managers in their real-world decisions.

References

Boulding, W., A. Kalra and R. Staelin (1999). *The quality double whammy*. Marketing Science, **18(4)**, 131–152.

Camerer, C. and E. Johnson (2004). *Thinking about attention in games: Backward and forward induction*. In: The Psychology of Economic Decisions, pp. 111–129. Oxford University Press.

Camerer, C. and D. Lovallo (1999). *Overconfidence and excess entry: An experimental approach*. The American Economic Review, **89(1)**, 428–443.

Camerer, C., T. H. Ho and J. K. Chong (2004). *A cognitive hierarchy model of games*. The Quarterly Journal of Economics, **119(3)**, 861–898.

Gabaix X. and D. Laibson (2000). *A boundedly rational decision algorithm*. American Economic Review, **90(2)**, 433–438.

Gigerenzer, G. and R. Selten (2001). *Bounded rationality: The adaptive toolbox*. MIT Press: Cambridge.

Hoffrage, U. (2004). *Cognitive Illusions: a handbook on fallacies and biases in thinking, judgement and memory*. Psychology Press: New York.

Hsu, T. A. (2002). *Behind Deep Blue*. Princeton University Press: Princeton.

Johnson, E. and J. Payne (1985). *Effort and accuracy in choice*. Management Science, **31(4)**, 395–414.

Johnson, E., C. Camerer, S. Sen and T. Rymon (2002). *Detecting failures of backward induction: Monitoring information search in sequential bargaining*. Journal of Economic Theory, **104(1)**, 16–47.

Narasimhan, C., C. He, E. Anderson, L. Brenner, P. Desai, D. Kuksov, P. Messinger, S. Moorthy, J. Nunes, Y. Rottenstreich, R. Staelin, G. Wu and Z. Zhang (2005). *Incorporating behavioral anomalies in strategic models*. Marketing Letters, **16(3)**, 361–373.

Newell, A. and H. A. Simon (1972). *Human problem solving*. Prentice-Hall: Englewood Cliffs, NJ.

Payne, J. W., J. R. Bettman, E. J. Johnson and M. F. Luce (1995). *An information processing perspective on choice*. In: Psychology of Learning and Motivation (J. Busemeyer, R. Hastie, and D. L. Medin, ed), Vol. 2, pp. 137–175. Academic Press: New York.

Payne, J. W., J. R. Bettman and M. F. Luce (1996). *When time is money: Decision behavior under opportunity-cost time pressure*. Organizational Behavior and Human Decision Processes, **66(2)**, 131–152.

Russell, S. and P. Norvig (2003). *Artificial Intelligence: A Modern Approach*. Pearson Education: Upper Saddle River, New Jersey.

Savage, L. (1954). *The foundations of statistics*. Wiley: New York.

Shugan, S. M. (1980). *The cost of thinking*. Journal of Consumer Research, **7(2)**, 99–111.

Stahl, D. O. (1993). *Evolution of smart players*. Games and Economic Behavior, **5**, 604–617.

Stahl, D. O. and P. W. Wilson (1995). *On player's models of other players: Theory and experimental evidence*. Games and Economic Behavior, **10(1)**, 218–254.