

On the Benefit of Nonlinear Control for Robust Logarithmic Growth: Coin Flipping Games as a Demonstration Case

Anton V. Proskurnikov and B. Ross Barmish

Abstract—The takeoff point for this paper is the voluminous body of literature addressing recursive betting games with expected logarithmic growth of wealth being the performance criterion. Whereas almost all existing papers involve use of linear feedback, the use of nonlinear control is conspicuously absent. This is epitomized by the large subset of this literature dealing with Kelly Betting. With this as the high-level motivation, we study the potential for use of nonlinear control in this framework. To this end, we consider a “demonstration case” which is one of the simplest scenarios encountered in this line of research: repeated flips of a biased coin with probability of heads p , and even-money payoff on each flip. First, we formulate a new robust nonlinear control problem which we believe is both simple to understand and apropos for dealing with concerns about distributional robustness; i.e., instead of assuming that p is perfectly known as in the case of the classical Kelly formulation, we begin with a bounding set $\mathcal{P} \subseteq [0, 1]$ for this probability. Then, we provide a theorem, our main result, which gives a closed-form description of the optimal robust nonlinear controller and a corollary which establishes that it robustly outperforms linear controllers such as those found in the literature. A second contribution of this paper bears upon the computability of our solution. For an n -flip game, whereas an admissible controller has $2^n - 1$ parameters, at the optimum only $O(n^2)$ of them turn out to be distinct.

I. INTRODUCTION

This paper addresses a large class of betting games described by discrete-time Markov processes. In this setting, the bettor begins with initial account value $V_0 > 0$. At each stage k , the control u_k , alternatively called the *betting strategy*, determines the size of the k -th wager. Then, over n steps, assuming independent identically distributed random variables X_k as the *returns* with a known probability distribution, the account value trajectory V_1, V_2, \dots, V_n , emanating from initial condition $V_0 > 0$, is obtained recursively by

$$V_{k+1} = V_k + u_k X_k.$$

Within this context, the literature most closely related to this paper concentrates on the design of a controller u maximizing the resulting *Expected Logarithmic Growth* (ELG)

$$\text{ELG}_u = \frac{1}{n} \mathbb{E} \log \left(\frac{V_n}{V_0} \right)$$

subject to *budget constraints* $|u_k| \leq V_k$ for $k = 0, 1, \dots, n-1$.

Perhaps, the most celebrated work along the lines above is the seminal paper by Kelly [1]; see also the early recognition

Anton V. Proskurnikov is associate professor with Department of Electronics and Telecommunications at Politecnico di Torino, Turin, Italy and B. Ross Barmish is an emeritus professor in ECE at University of Wisconsin, Madison, and CEO of Robust Trading Solutions, LLC.

anton.p.1982@ieee.org, bob.barmish@gmail.com

of the power of the ELG approach in [2]–[5]. Over the decades to follow, we see a voluminous body of literature, comprised of hundreds of papers, dealing with applications, extensions, and generalizations of Kelly’s result in various directions. We also see many papers providing rationale for the use of the logarithmic growth criterion versus other performance metrics. A selection of highlights from this work includes the detailed coverage of these topics in textbooks such as [6] and [7] and the extensive collection of papers in [8]. It is also important to point out that this body of literature being cited includes major results on various properties of the ELG-maximizing controller over and above optimal logarithmic growth. That is, many authors provide results bearing on the “asymptotic superiority” of the ELG maximizer and cover other topics such as the relaxation the i.i.d. assumption on the X_k . In this regard, some good starting points for the uninitiated reader are [9]–[11]. Finally, we draw attention to the doctoral dissertation of Hsieh [12] which includes not only a comprehensive review of the earlier literature but also details and citations of his contributions and those of others over the preceding years.

To complete this brief perspective of the literature related to this paper, it is also important to mention the body of work dealing with “distributional robustness” issues arising in stochastic optimization; see [13] and [14] where this terminology is introduced, the development of the theory in [15], the important 2016 paper [16], dealing specifically with distributional robustness in an ELG context and the more recent ongoing work [17] for along these lines.

Given the research setting above, the primary motivation for this paper is the fact that in the existing ELG literature, a problem formulation as one of nonlinear control is conspicuously absent; i.e., only linear control is considered. Whereas it is arguable, based on some of the results in existing work, that a nonlinear control cannot outperform a linear feedback when the probability distribution for the X_k is perfectly known, our main contention in this paper is that the same does not hold true when uncertainty in the underlying probability distributions is in play. Said another way, our main results provide compelling evidence that there are a large number of scenarios, involving distributional robustness considerations for which a nonlinear controller can outperform the “best” linear controller; e.g., see [16] and [18]. To this end, our analysis to follow demonstrates the potential of robust nonlinear control by considering one of the simplest possible ELG scenarios: a coin-flipping game with uncertainty in the probability of heads p . Instead of

taking p to be perfectly known as in the case of the classical Kelly formulation, we begin with a bounding set $\mathcal{P} \subseteq [0, 1]$ for this probability. In this setting, the main result in this paper is a theorem which provides a complete closed-form solution of an optimal robust nonlinear control problem. As a corollary, we prove that our nonlinear controller robustly outperforms any linear controller.

A second contribution of this paper bears on the computability of our new solution. Whereas an admissible controller has $2^n - 1$ design parameters associated with the sample path points for an n -flip game, surprisingly, at the optimum, many of them turn out to be the same with the resulting number of “free parameters” being of $O(n^2)$. Finally, we provide some initial illustrations bearing on the “cost of imprecision” in our knowledge of the probability p . To this end, some comparisons are made between the ELG associated with the optimal robust nonlinear controller and the so-called *perfect-information Kelly optimum*.

Our Demonstration Case: Coin-Flipping

To demonstrate the potential for the use of nonlinear control, we analyze one of the simplest and most fundamental problems in the logarithmic growth literature: Making bets on n consecutive flips of a biased coin with probability of heads being p and even-money payoff. This initial assumption on the payoffs is made for simplicity and brevity of the exposition; see Section VII for a generalization to the unequal payoff case. Our simplified framework enables us to explain the key ideas behind our new nonlinear control formulation without being encumbered by additional details.

Indeed, to demonstrate the potential for consideration of nonlinear controllers in future research, we begin with the widely celebrated betting scheme of Kelly [1]. That is, using the notation above, at stage k , the controller generates the bet size as a linear feedback. That is, $u_k = KV_k$ with the understanding that $u_k > 0$ and $u_k < 0$ corresponds to bets on heads and tails respectively. Consistent with this, even money payoffs are $X_k = 1$ for heads and $X_k = -1$ for tails.

Then, with probability of heads p assumed to be perfectly known and budget constraint $|u(k)| \leq V_k$ imposed, a straightforward calculation leads to Kelly’s optimal ELG maximizing feedback gain $K^* = 2p - 1$. It is important to note that, this simple linear controller proves to be optimal in many cases other than the simple scenario described above; e.g., in the widely cited 1971 paper by Hakansson [9] (see also [6]), the bet size u_k may depend on the entirety of the past history $V_0, V_1, \dots, V_{k-1}, V_k$. However, as indicated earlier, our goal is to demonstrate the importance of nonlinear control when p is imperfectly known.

II. ADMISSIBLE NONLINEAR CONTROLLERS AND RESULTING EXPECTED LOGARITHMIC GROWTH

For the coin-flipping game at hand with sample path space

$$\mathcal{X} \doteq \{-1, 1\}^n,$$

a mapping $K : \mathcal{X} \rightarrow \mathbb{R}^n$ is said to define an *admissible nonlinear controller* if the following conditions are satisfied:

First, mapping K is *causal*; that is, K_0 is a constant, and, given any sample path $X = (X_0, X_1, \dots, X_{n-1}) \in \mathcal{X}$, for $k = 1, 2, \dots, n - 1$, the controller’s k -th component $K_k(X)$ depends only on X_0, X_1, \dots, X_{k-1} , with the resulting bet size at stage k given by $u_k(K, X) = K_k(X)V_k(K, X)$. Second, controller K should satisfy the *budget constraint* $|u_k(K, X)| \leq V_k(K, X)$, i.e., $|K_k(X)| \leq 1$ for all $k = 0, \dots, n - 1$. In the sequel, we denote the set of all admissible controllers by \mathcal{K} .

It is important to note that \mathcal{K} includes linear controllers as special case. More generally, members of \mathcal{K} can be highly nonlinear functions of X . Accordingly, whenever appropriate, $K_k(X)$ is referred to as a *nonlinear feedback gain*. Now, along sample path $X \in \mathcal{X}$, the resulting account value is obtained recursively as $V_0(K, X) = V_0$, and

$$V_{k+1}(K, X) = (1 + K_k(X)X_k)V_k(K, X), \quad k = 0, 1, \dots$$

In Figure 1, the binary tree associated with the state transitions above are shown. We also draw attention to the color scheme used for the nodes: At any given stage k , two nodes with the same color represent sample paths with the same number of heads over the prior stages $0, 1, 2, \dots, k - 1$. As seen in the main result to follow, at such nodes, the *optimal* robust nonlinear control, has the same nonlinear gain $K_k(X)$. Now, continuing with the analysis, the final logarithmic growth is

$$\frac{1}{n} \log \left(\frac{V_n(K, X)}{V_0} \right) = \frac{1}{n} \sum_{k=0}^{n-1} \log((1 + K_k(X)X_k)).$$

We now turn our attention to the starting point for much of the analysis to follow: the simple formula for the Expected Logarithmic Growth (ELG) as a function of the controller nonlinear gains $K \in \mathcal{K}$ and the probability of heads p . We first find the probability of a sample path $X \in \mathcal{X}$, given by

$$P(X) \doteq p^{n_h(X)}(1 - p)^{n - n_h(X)}$$

where $n_h(X) \doteq \#\{i = 0, \dots, n - 1 : X_i = 1\}$ is the number of heads. The ELG as a function of K and p is found as¹

$$\text{ELG}_K(p) = \frac{1}{n} \sum_{X \in \mathcal{X}} P(X) \sum_{k=0}^{n-1} \log(1 + K_k(X)X_k).$$

III. ROBUSTNESS FORMULATION

Per earlier discussion, we now formulate a Robust Expected Logarithmic Growth problem involving uncertainty in the probability of heads. To this end, let $\mathcal{P} \subseteq [0, 1]$ denote a Lebesgue measurable set of the possible values for p against which we seek robustness, and, let $\mu(\mathcal{P})$ be its corresponding Lebesgue measure. For example, if the only a priori information we have about the probability of heads are bounds

$$0 \leq p_{min} \leq p \leq p_{max} \leq 1,$$

¹Formally, it is possible that $V_n(K, X) = 0$ for some $X \in \mathcal{X}$. If $P(X) > 0$ for such a sample path, then $\text{ELG}_K(p) = -\infty$; otherwise, the resulting summand is neglected by using the convention $0 \cdot (-\infty) = 0$.




Fig. 1. Causal Controllers and Random Walks on a Binary Tree

then with $\mathcal{P} = [p_{min}, p_{max}]$, we have $\mu(\mathcal{P}) = p_{max} - p_{min}$. In the sequel, to avoid trivialities, we assume $\mu(\mathcal{P}) > 0$. Unlike Kelly's perfect-information scenario, we cannot take $K \in \mathcal{K}$ depending on the unknown probability of heads p ; the controller $K = K_{\mathcal{P}}$ should be determined by the known set \mathcal{P} .

Comparison With Kelly's Perfect-Information Optimum

To assess the robustness of any particular controller, we compare it's expected logarithmic growth, as a function of $p \in \mathcal{P}$, with that of Kelly's *perfect-information ELG optimum* $K_p^* \doteq 2p - 1$ described in the Introduction. In this regard, we view K_p^* as a member the admissible set \mathcal{K} , corresponding to optimal performance level

$$\text{ELG}^*(p) = p \log(2p) + (1 - p) \log(2(1 - p)).$$

This quantity, serves as our "gold standard" against which we assess the robust performance of controllers with *imperfect information*. We first observe that the inequality $\text{ELG}^*(p) \geq \text{ELG}_K(p)$ must hold for all $p \in \mathcal{P}$ and $K \in \mathcal{K}$. Hence, the associated error integral

$$\text{Err}(K) \doteq \int_{p \in \mathcal{P}} (\text{ELG}^*(p) - \text{ELG}_K(p)) dp$$

is minimized by maximizing the Integral Expected Logarithmic Growth (IELG), that is, the function

$$\text{IELG}_K \doteq \int_{p \in \mathcal{P}} \text{ELG}_K(p) dp$$

over all admissible controllers $K \in \mathcal{K}$. By viewing p as a random variable uniformly distributed over \mathcal{P} , the IELG is in fact proportional to the *expectation* of random variable ELG_K :

$$\mathbb{E}(\text{ELG}_K) = \frac{1}{\mu(\mathcal{P})} \int_{p \in \mathcal{P}} \text{ELG}_K(p) dp = \frac{\text{IELG}_K}{\mu(\mathcal{P})}.$$

It is noted that use of a uniform distribution, implicit in the integral above, is being used solely for the sake of simplicity of the exposition. As discussed in Section VII, the analysis to follow is easily modified to address more general distributions on \mathcal{P} . In the theorem to follow in Section VI, it is seen that element $K^* \in \mathcal{K}$ maximizing the IELG exists, is unique, and is given by a simple and efficient formula.

IV. THE SUBCLASS OF STATIC LINEAR CONTROLLERS

By way of preliminaries, we say that an admissible $K \in \mathcal{K}$ defines an *admissible static linear controller* if there exists a constant $K_0 \in [-1, 1]$ such that $K(X) \equiv K_0$ for all $X \in \mathcal{X}$. In order to study the robust performance over this subclass of controllers, we work with the function

$$\text{IELG}_{K_0} = \int_{p \in \mathcal{P}} (p \log(1 + K_0) + (1 - p) \log(1 - K_0)) dp.$$

Lemma: *The static feedback gain maximizing IELG_{K_0} subject to the constraint $K_0 \in [-1, 1]$ is unique and given by*

$$K_0^* = 2\bar{p} - 1; \quad \bar{p} \doteq \frac{1}{\mu(\mathcal{P})} \int_{p \in \mathcal{P}} p dp.$$

Proof: A straightforward computation indicates that for $K_0 \in (-1, 1)$, function $f(K_0) \doteq \text{IELG}_{K_0}$ has the first derivative

$$f'(K_0) = \frac{2 \int p dp - \mu(\mathcal{P}) - K_0 \mu(\mathcal{P})}{1 - K_0^2} = \frac{\mu(\mathcal{P})(2\bar{p} - 1 - K_0)}{1 - K_0^2}$$

whereas $f(\pm 1) = -\infty$. Since this derivative is positive when $K_0 < K_0^*$, negative when $K_0 > K_0^*$ and zero at $K_0 = K_0^*$, it follows that K_0^* maximizes IELG_{K_0} over $K_0 \in [-1, 1]$. \square

Observations and Important Special Cases: Notice that the optimal robust static linear control coincides with the Kelly's formula K_p^* for the perfect-information case with probability of heads p being the *centroid* \bar{p} of uncertainty set \mathcal{P} . It is also interesting that we can obtain Kelly's formula K_p^* as a special case of our robustness analysis by relaxing the positive measure assumption on \mathcal{P} and using a limiting argument. That is, suppose $p \in (0, 1)$ and $\mathcal{P} = [p, p + \delta]$ with parameter δ satisfying $0 < \delta \leq 1 - p$. Then, taking the limit as $\delta \rightarrow 0$, it is straightforward to verify that the optimal robust linear feedback gain reduces to $K_0^* = 2p - 1$. Next, we consider the case when \mathcal{P} is a union positive-length disjoint intervals

$$\mathcal{P} = \bigcup_{i=1}^m [p_{min,i}, p_{max,i}].$$

Applying the lemma, a straightforward calculation leads to

$$K_0^* = \frac{\sum_{i=1}^m (p_{max,i}^2 - p_{min,i}^2)}{\sum_{i=1}^m (p_{max,i} - p_{min,i})} - 1$$

which specializes further: If all differences $p_{max,i} - p_{min,i}$ are the same, then, the formula above reduces to

$$K_0^* = \frac{1}{m} \sum_{i=1}^m (p_{min,i} + p_{max,i}) - 1$$

For the single-interval case $\mathcal{P} = [p_{min}, p_{max}]$, we have

$$K_0^* = p_{min} + p_{max} - 1.$$

V. EXAMPLES: OPTIMAL NONLINEAR CONTROLLER

To motivate the key ideas underlying the general result in the theorem to follow, we calculate the optimal robust nonlinear control for the simpler special cases $n = 2$ and $n = 3$. As seen below, each of these optima can be found in a very simple manner. That is, each of the desired nonlinear

feedback gains comprising the optimum $K^* \in \mathcal{K}$ is found via single-variable maximization whose solution admits a closed form. The simple argument used in these examples is also important in the proof for the general case of $n > 2$ in the theorem to follow.

Example 1: Beginning with the $n = 2$ and $\mathcal{P} = [0, 1]$, to simplify calculations, we first represent the nonlinear controller components $K_k(X)$ employing the shorthand notation $a = K_0(X)$ for all $X \in \mathcal{X}$, $b = K_1(1, 1) = K_1(1, -1)$ and $c = K_1(-1, 1) = K(-1, -1)$, we first calculate

$$\begin{aligned} \text{ELG}_K(p) &= \\ &= \frac{1}{2} [p^2 \log(1+a)(1+b) + p(1-p) \log(1+a)(1-b) + \\ & p(1-p) \log(1-a)(1+c) + (1-p)^2 \log(1-a)(1-c)]. \end{aligned}$$

Then, upon expanding the logarithms above and integrating with respect to $p \in \mathcal{P}$, we obtain the IELG function

$$\begin{aligned} \text{IELG}_K &= \frac{1}{2} \left[\frac{1}{2} \log(1+a) + \frac{1}{2} \log(1-a) \right. \\ & \quad \left. + \frac{1}{3} \log(1+b) + \frac{1}{6} \log(1-b) \right. \\ & \quad \left. + \frac{1}{6} \log(1+c) + \frac{1}{3} \log(1-c) \right]. \end{aligned}$$

Next, we note that the desired optimization with respect (a, b, c) above can be solved by maximizing three separate single-variable strictly concave functions; each of functions is of the form $f(x) = \alpha \log(1+x) + \beta \log(1-x)$ with unconstrained maximum, obtained by setting the derivative to zero, given by $x^* = (\alpha - \beta)/(\alpha + \beta)$. Then, we simply observe that the budget-constrained optimum is also x^* because in all three cases above, α and β are such that $|x^*| \leq 1$. Based on these considerations, we immediately arrive at a unique maximizing solution $a = 0; b = \frac{1}{3}; c = -\frac{1}{3}$. In other words, it is optimal to skip the first bet, and bet $1/3$ of the account value on heads if the first toss comes up to be heads ($X_0 = 1$), and bet $1/3$ of the account value on tails if the first toss comes up to be tails ($X_0 = -1$). Furthermore, via a straightforward substitution we obtain $\text{ELG}_{K^*} = \frac{1}{3} \log \frac{32}{27} \approx 0.0566$. Note that the optimal static *linear* controller from Lemma 1 is degenerate: $K_0^* = 0$ and $\text{ELG}_{K_0^*} = 0$.

Example 2: Now proceeding to the analysis for $n = 3$, we consider the case $\mathcal{P} = [0.25, 0.95]$ and again represent the nonlinear control gains using a shorthand notation which takes causality into account: For each $K_k(X)$, only the first k components of its argument X are indicated and we take

$$\begin{aligned} a &= K_0; \quad b = K_1(1); \quad c = K_2(1, 1); \quad d = K_2(1, -1); \\ e &= K_1(-1); \quad f = K_2(-1, 1); \quad g = K_2(-1, -1). \end{aligned}$$

Now, we repeat the straightforward sum-of-logarithms computation as in Example 1 to arrive at

$$\begin{aligned} \text{IELG}_K &= \frac{1}{3} [0.09333 \log(1-a) + 0.1400 \log(1+a) \\ & \quad + 0.04647 \log(1-b) + 0.09353 \log(1+b) \\ & \quad + 0.02598 \log(1-c) + 0.06755 \log(1+c) \\ & \quad + 0.02049 \log(1-d) + 0.02598 \log(1+d) \\ & \quad + 0.04686 \log(1-e) + 0.04647 \log(1+e) \\ & \quad + 0.02049 \log(1-f) + 0.02598 \log(1+f) \\ & \quad + 0.02637 \log(1-g) + 0.02049 \log(1+g)]. \end{aligned}$$

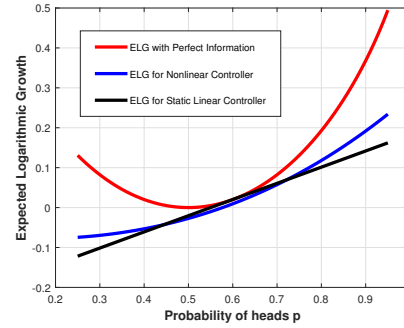


Fig. 2. Robust Performance Plots for Comparison Purposes

Now, proceeding again as in Example 1, we maximize separately with respect to each of the seven control parameters and obtain unique optimum $a = 0.2, b = 0.3361, c = 0.4445, d = 0.118, e = -0.004167, f = 0.118, g = -0.007429$. For performance comparison, we calculate the optimal static gain $K_0^* = p_{min} + p_{max} - 1 = 0.2$.

It is interesting to compare the robust performance for the optimal nonlinear controller $\text{ELG}_{K^*}(p)$ with that of the optimal static linear controller $\text{ELG}_{K_0^*}(p)$, as functions of the probability $p \in \mathcal{P}$ and benchmarked with Kelly's perfect-information optimum $\text{ELG}^*(p)$ as the best-possible upper bound. In Figure 2, where these quantities are plotted, the following is noted: While the optimal nonlinear controller leads to a larger IELG (the area under the curve), at some specific values of p the static linear controller outperforms it in terms of $\text{ELG}(p)$. This is not surprising, because at the midpoint $p = 0.6$ of interval for \mathcal{P} , the optimal static gain K_0^* , in fact, coincides with the Kelly gain $K^*(p)$.

VI. MAIN RESULT ON ROBUST OPTIMAL CONTROL

The theorem below establishes the existence and uniqueness of an optimal robust nonlinear controller $K^* \in \mathcal{K}$ and characterizes it with an explicit formula. Although there are $2^n - 1$ nonlinear controllers gains associated with the nodes, the theorem tells us that at each stage k , there are only $k + 1$ possible values of for the optimal robust nonlinear gain $K_k^*(X)$. Summing up these numbers across all stages, we see that the total number of nonlinear gains to be calculated is $n(n+1)/2$; i.e., the computational burden of finding all gains $K_k(X)$ increases quadratically in n rather than exponentially. On the other hand, if the time between consecutive coin flips is suitably large, the controller in the theorem below can be implemented "dynamically" with no need to pre-compute the optimal nonlinear gains $K_k^*(\bar{X})$.

Theorem: *The integral expected logarithmic growth IELG_K defined over on the set of admissible controllers $K \in \mathcal{K}$, has a unique maximizer K^* , whose nonlinear control gain at stage k for a sample path $\bar{X} \in \mathcal{X}$ is*

$$K_k^*(\bar{X}) = \frac{\int_{p \in \mathcal{P}} p^{q_k} (1-p)^{k-q_k} (2p-1) dp}{\int_{p \in \mathcal{P}} p^{q_k} (1-p)^{k-q_k} dp},$$

where $q_k = q_k(\bar{X}) = \#\{i = 0, \dots, k-1 : \bar{X}_i = 1\} \leq k$ is the number of heads occurring over the first k coin flips².

Proof: In the arguments to follow, sample path $\bar{X} \in \mathcal{X}$ and stage number $k \in \{0, 1, \dots, n-1\}$ are assumed to be fixed, and we let $\mathcal{X}_k(\bar{X}) \doteq \mathcal{X}_k^+(\bar{X}) \cup \mathcal{X}_k^-(\bar{X})$, where³

$$\mathcal{X}_k^+(\bar{X}) \doteq \{X \in \mathcal{X} : X_i = \bar{X}_i \forall i = 0, \dots, k-1; X_k = 1\},$$

$$\mathcal{X}_k^-(\bar{X}) \doteq \{X \in \mathcal{X} : X_i = \bar{X}_i \forall i = 0, \dots, k-1; X_k = -1\},$$

the expected logarithmic growth function can be written as

$$\begin{aligned} \text{ELG}_K(p) = & \frac{1}{n} \sum_{X \in \mathcal{X}_k^+(\bar{X})} P(X) \sum_{i=0}^{n-1} \log(1 + K_i(X)X_i) \\ & + \frac{1}{n} \sum_{X \in \mathcal{X}_k^-(\bar{X})} P(X) \sum_{i=0}^{n-1} \log(1 + K_i(X)X_i) \\ & + \frac{1}{n} \sum_{X \notin \mathcal{X}_k(\bar{X})} P(X) \sum_{i=0}^{n-1} \log(1 + K_i(X)X_i). \end{aligned}$$

Next we note that all of the terms above involving $K_k(\bar{X})$ can be isolated by setting $i = k$ in the first two terms above. As far as the third term is concerned, it is independent of $K_k(\bar{X})$ because admissibility of the controller forces $K_k(X) = K_k(\bar{X})$ for all $X \in \mathcal{X}_k(\bar{X})$. Now integrating $\text{ELG}_K(p)$ over $p \in \mathcal{P}$, it is straightforward to see that maximization of IELG_K over $K \in \mathcal{K}$ reduces to maximization of the single-variable function

$$g_{k,\bar{X}}(K_k) \doteq \alpha_k(\bar{X}) \log(1 + K_k) + \beta_k(\bar{X}) \log(1 - K_k)$$

over the interval $K_k \in [-1, 1]$. Here α_k, β_k are defined as

$$\alpha_k(\bar{X}) \doteq \int_{p \in \mathcal{P}} \sum_{X \in \mathcal{X}_k^+(\bar{X})} P(X) dp = \int_{p \in \mathcal{P}} \mathbb{P}(X \in \mathcal{X}_k^+(\bar{X})) dp,$$

$$\beta_k(\bar{X}) \doteq \int_{p \in \mathcal{P}} \sum_{X \in \mathcal{X}_k^-(\bar{X})} P(X) dp = \int_{p \in \mathcal{P}} \mathbb{P}(X \in \mathcal{X}_k^-(\bar{X})) dp.$$

To find α_k , notice that event $X \in \mathcal{X}_k^+$ is the intersection of $k+1$ events $X_i = \bar{X}_i$ (where $i = 0, \dots, k-1$) and $X_k = 1$; these events are mutually independent. Among these events, there are $1 + q_k(\bar{X})$ events of type $X_j = 1$ and probability p and $k - q_k(\bar{X})$ events of type $X_j = -1$ and probability $1 - p$. Multiplying these probabilities, one has

$$\alpha_k(\bar{X}) = \int_{p \in \mathcal{P}} p^{q_k(\bar{X})+1} (1-p)^{k-q_k(\bar{X})} dp.$$

A very similar argument yields in the expression for β_k ; i.e.,

$$\beta_k(\bar{X}) = \int_{p \in \mathcal{P}} p^{q_k(\bar{X})} (1-p)^{k+1-q_k(\bar{X})} dp.$$

²By definition, $q_0(\bar{X}) = 0$, i.e., the optimal control gain $K_0^*(X) \equiv K_0^*$ at stage $k = 0$ coincides with the optimal linear controller from Lemma 1.

³For $k = 0$, sets \mathcal{X}_k^+ and \mathcal{X}_k^- consist of all sample paths starting from $X_0 = 1$ and $X_0 = -1$, respectively.

Since $g_{k,\bar{X}}(K_k)$ is strictly concave on interval $K_k \in [-1, 1]$ and $g_{k,\bar{X}}(\pm 1) = -\infty$, the optimum is found by setting the derivative with respect to K_k to zero; i.e., we obtain

$$\begin{aligned} K_k^*(\bar{X}) &= \frac{\alpha_k(\bar{X}) - \beta_k(\bar{X})}{\alpha_k(\bar{X}) + \beta_k(\bar{X})} = \\ &= \frac{\int_{p \in \mathcal{P}} p^{q_k(\bar{X})} (1-p)^{k-q_k(\bar{X})} (2p-1) dp}{\int_{p \in \mathcal{P}} p^{q_k(\bar{X})} (1-p)^{k-q_k(\bar{X})} dp}, \end{aligned}$$

which satisfies $|K_k^*(\bar{X})| \leq 1$ as required. \square

Nonlinear Versus Linear Control

We are now prepared to address one of our main contentions articulated the title and abstract. That is, except for the trivial case of single-flip game ($n = 1$), we establish, as a corollary of the theorem, that the optimal nonlinear controller $K^* \in \mathcal{K}$ robustly outperforms the optimal static linear feedback K_0^* .

Corollary: For $n > 1$ steps, one has $\text{IELG}_{K^*} > \text{IELG}_{K_0^*}$.

To facilitate the proof, we provide a preliminary lemma.

Preliminary Lemma: Let $\mathcal{P} \subseteq [0, 1]$ be a Lebesgue measurable set with $\mu(\mathcal{P}) > 0$. Then, for all $n > 1$, one has

$$\int_{p \in \mathcal{P}} p^{n-1} dp \int_{p \in \mathcal{P}} p dp < \mu(\mathcal{P}) \int_{p \in \mathcal{P}} p^n dp.$$

Proof. Denoting $I_k \doteq \int_{p \in \mathcal{P}} p^k dp$, we need to prove that $I_{n-1}I_1 < I_0I_n$. Applying Hölder's inequality to pairs of functions $f(p) = p, g(p) = 1$ and $g(p) = 1, h(p) = p^{n-1}$ and conjugate exponents n and $m = n/(n-1)$, one obtains

$$I_1 = \|fg\|_{L_1(\mathcal{P})} < \|f\|_{L_n(\mathcal{P})} \|g\|_{L_m(\mathcal{P})} = I_n^{1/n} I_0^{1/m},$$

$$I_{n-1} = \|gh\|_{L_1(\mathcal{P})} < \|g\|_{L_n(\mathcal{P})} \|h\|_{L_m(\mathcal{P})} = I_0^{1/n} I_n^{1/m}.$$

Multiplying these two inequalities, it follows that $I_1I_{n-1} < I_n^{1/n} I_0^{1/m} I_0^{1/n} I_n^{1/m} = I_n I_0$ finishing the proof. \square

Proof of Corollary: Recalling the theorem, K^* is the unique maximizer of IELG_K over all admissible nonlinear gains $K \in \mathcal{K}$. Since the optimal linear static gain K_0^* is admissible, it suffices to show that $K_k^*(X) \neq K_0^*$ for at least one $k \in \{0, 1, \dots, n-1\}$ and at least one sample path $X \in \mathcal{X}$. Indeed, considering the distinguished sample path corresponding to all heads; i.e., $X_0 = \dots = X_{n-1} = 1$, we first note that $q_{n-1}(X) = n$. Recall that $K_0^* = 2\bar{p} - 1$, where \bar{p} is the centroid of \mathcal{P} . Now applying the theorem, we have

$$K_{n-1}^*(X) = 2\hat{p}_n - 1; \quad \hat{p}_n \doteq \frac{\int_{p \in \mathcal{P}} p^n dp}{\int_{p \in \mathcal{P}} p^{n-1} dp} > \bar{p}$$

with the latter inequality implied by our Preliminary Lemma. Thus, $K_{n-1}^*(X) > K_0^*$, and hence K_0^* is non-optimal. \square

VII. DISCUSSION OF TWO GENERALIZATIONS

As previously mentioned, the assumption of equal payoffs for heads and tails was made solely for simplicity of exposition and brevity of the presentation. We now sketch the key ideas indicating how the robustly optimal nonlinear gains $K_k^*(X)$ are obtained for the unequal payoff case with $X_k = a$ for heads and $X_k = b$ for tails at stage k ;

it is assumed that $b < 0 < a$. Let $q_{a,k} \doteq q_{a,k}(X)$ be the number of heads seen on the first k flips and

$$\alpha_{a,k} \doteq \int_{\mathcal{P}} p^{q_{a,k}+1} (1-p)^{k-q_{a,k}} dp,$$

$$\beta_{a,k} \doteq \int_{\mathcal{P}} p^{q_{a,k}} (1-p)^{k-q_{a,k}+1} dp.$$

Then, to get the robustly optimal gain $K_k^* = K_k^*(X)$, we form the *strictly concave* scalar function

$$g_k(K_k) \doteq \alpha_{a,k} \log(1 + K_k a) + \beta_{a,k} \log(1 + K_k b)$$

to be maximized subject to budget constraint $-1 \leq K \leq 1$ and the requirement $V_{k+1} \geq 0$ associated with both well-definedness of the logarithms above and bankruptcy considerations. Accordingly, with

$$m \doteq \max \left\{ -1, -\frac{1}{a} \right\}; \quad M \doteq \min \left\{ 1, \frac{1}{|b|} \right\}$$

we obtain the optimal nonlinear gain K_k^* as the unique maximizer of $g_k(K_k)$ on the interval $K_k \in [m, M]$.

Finally, as has been already mentioned earlier in the paper, our results retain their validity if one replaces the uniform distribution on \mathcal{P} by a positive finite measure $\hat{\mu}$ (defined, at least, on Borel subsets of \mathcal{P}) and redefining the IELG as

$$\text{IELG}_K = \int_{\mathcal{P}} \text{ELG}_K(p) \hat{\mu}(dp).$$

Our Lemma, Theorem and Corollary then remain valid, replacing dp in all integrals by $\hat{\mu}(dp)$ and $\mu(P)$ by $\hat{\mu}(P)$.

VIII. CONCLUSION AND FUTURE RESEARCH

In this paper, our main objective was to demonstrate that nonlinear control has an important role to play in large classes of betting games dealing with Expected Logarithmic Growth. To this end, we considered a simple coin-flipping game as a demonstration case to convey our main ideas. Whereas a static linear control is “unbeatable” with a perfectly known probability of heads p , this does not hold true when robustness with respect to variations in p is of concern. For this situation, we showed that the optimal controller with its nonlinear gains K^* robustly outperforms the optimal static linear controller with its gain K_0^* .

Perhaps the main implication of our results is that future study of nonlinear control with more general problem formulations, relevant for the field of finance, is likely to bear fruit. By way of future research, in addition to the generalizations sketched in Section VII, we believe that it should be possible to address the case when the returns X_k take on multiple or even a continuum of values governed by rather general probability distributions. Another possible generalization begins with “vector sample paths” X in lieu of the scalar ones considered here. Such a formulation can be viewed in a robust portfolio balancing context with results along the lines serving as a stepping stone to applications such as algorithmic stock trading in financial markets. In bringing such results from theory to practice, it would be important to add terms to the account value dynamics for V_k

to include consideration of factors such as the risk-free and margin interest rates, leverage and transaction costs.

Finally, we mention one additional continuation of this research which is motivated by the following observation: An adaptive controller aimed at maximizing expected logarithmic growth, say along the lines of those given in recent papers such as [19] and [20], should rightfully be viewed as a member of our admissible control set \mathcal{K} . Accordingly, our plan for future research involves exploring the connection between results in adaptive and nonlinear control which have traditionally been viewed as rather separate areas.

REFERENCES

- [1] J. L. Kelly, “A New Interpretation of Information Rate,” *Bell System Technical Journal*, vol. 35.4, pp. 917–926, 1956.
- [2] L. Breiman, “Optimal Gambling Systems for Favourable Games,” *Fourth Berkeley Symposium on Mathematical Statistics and Probability*, University of California Press, pp. 65–78, 1961.
- [3] E. O. Thorp, “Fortune’s Formula: The Game of Blackjack,” *American Mathematical Society*, 1961.
- [4] E. O. Thorp, “Optimal Gambling Systems for Favorable Games,” *Review of the International Statistical Institute*, vol. 37, pp. 273–293, 1969.
- [5] T. M. Cover, “Algorithm for Maximizing Expected Log Investment Return,” *IEEE Transactions on Information Theory*, IT-30, pp. 369–373, 1984.
- [6] T. M. Cover and T. J. Thomas, *Elements of Information Theory*, Wiley Series in Telecommunications, 1991.
- [7] D. G. Luenberger, *Investment Science*, Oxford University Press, 1998.
- [8] L. C. MacLean, E. O. Thorp, and W. T. Ziemba, *The Kelly Capital Growth Investment Criterion: Theory and Practice*. World Scientific, 2011.
- [9] N. H. Hakansson, “On Optimal Myopic Portfolio Policies With and Without Serial Correlation of Yields,” *Journal of Business*, vol. 44, pp. 324–334, 1971.
- [10] P. H. Algoet and T. M. Cover, “Asymptotic Optimality and Asymptotic Equipartition Properties of Log-Optimum Investment,” *The Annals of Probability*, vol. 16, pp. 876–898, 1988.
- [11] J. D. O’Brien, K. Burke, M. E. Burke, and B. R. Barmish, “A Generalization of the Classical Kelly Betting Formula to the Case of Temporal Correlation,” *IEEE Control Systems Letters*, vol.5, pp. 623–628, 2021.
- [12] C. H. Hsieh, “Contributions to the Theory of Kelly Betting with Applications to Stock Trading: A Control-Theoretic Approach,” Doctoral Dissertation, ECE Department, University of Wisconsin, 2019.
- [13] B. R. Barmish and P. S. Shcherbakov, “Distributionally Robust Least Squares,” *Proceedings of SPAS’99*, St. Petersburg, Russia, 1999.
- [14] C. M. Lagoa and B.R. Barmish, “Distributionally Robust Monte Carlo Simulation: A Tutorial Survey,” *Proceedings of the IFAC World Congress*, vol. 35, pp. 151–162, Barcelona Spain, 2002.
- [15] E. Y. Delage and Y. Ye, “Distributionally Robust Optimization Under Moment Uncertainty with Application to Data-Driven Problems,” *Operations Research*, vol. 58, pp. 595–612, 2010.
- [16] N. Rujeeapaiboon, D. Kuhn and W. Wiesemann, “Robust Growth-Optimal Portfolios,” *Management Science*, vol. 62, pp. 2090–2109, 2016.
- [17] J. Y. Li, “Wasserstein-Kelly Portfolios: A Robust Data-Driven Solution to Optimize Portfolio Growth,” <https://arxiv.org/abs/2302.13979>, 2023.
- [18] Q. Sun and S. Boyd, “Distributional Robust Kelly Strategy: Optimal Strategy under Uncertainty in the Long-Run,” arXiv preprint: 1812.10371, 2018.
- [19] A. Despons, L. Peliti and D. Lacoste, “Adaptive Strategies in Kelly’s Races Model,” *Journal of Statistical Mechanics: Theory and Experiment*, pp. 1–17, 2022.
- [20] F. Dettu, F. Abbracciavento and S. Formentin, “Kelly-Based Stock Trading via Feedback Control,” *Proceedings of the IEEE Conference on Decision and Control*, pp. 5574–5579, Cancun, Mexico, 2022.