

The Economist's View of Combinatorial Games

ELWYN BERLEKAMP

ABSTRACT. We introduce two equivalent methodologies for defining and computing a position's mean (value of playing Black rather than White) and temperature (value of next move). Both methodologies apply in more generality than the classical one. The first, following the notion of a free market, relies on the transfer of a "tax" between players, determined by continuous competitive auctions. The second relies on a *generalized thermograph*, which reduces to the classical thermograph when the game is loop-free.

When a sum of games is played optimally according the economic rules described, the mean (which is additive) and the temperature determine the final score precisely.

This framework extends and refines several classical notions. Thus, finite games that are numbers in Conway's sense are now seen to have *negative* natural temperatures. All games can now be viewed as terminating naturally with integer scores when the temperature reaches -1 .

Introduction

At every position of a game such as Go or Domineering, there are two very important questions:

Who is ahead, and by how much?

How big is the next move?

Following Conway [1976], classical abstract combinatorial game theorists answer these questions with a *value* and an *incentive*, as discussed in *Winning Ways* [Berlekamp et al. 1982]. These answers are precisely correct when the objective of the game is to get the last legal move. Values and incentives are themselves games, and can quickly become complicated.

Our ideal economist takes a different view. Following Hanner [1959] and Milnor [1953], he views the game as a contest to accumulate points, which can eventually be converted into cash. Our modern economist monetarizes the answers to our opening two questions into prices—real numbers that can be determined by competitive, free-market auctions. Specifically, after two gurus have completed their studies of a position, the economist might ask each to submit a sealed bid, representing the amount the guru is willing to pay in order to play Black. The bid may be any real number, including a negative one (if the position favors White). Assuming the bids differ, the referee computes their average μ . The

player with the higher bid plays Left (= Black), in return for which he pays μ dollars to his opponent. The player with the lower bid plays Right (= White). Each player is necessarily happy with this assignment, since it is no worse for him than the bid he submitted.

For mnemonic purposes, the reader should think of μ as the *market value*. Later, we shall see that it can also be interpreted as the *mean value*, or, in thermography, as the *mast value*.

Our ideal economist also uses competitive auctions to determine the sizes of the moves. Throughout the game, the economist requires that each move made by either player must be accompanied by a fee or “tax”, which the mover must pay to his opponent. At the beginning of the game, the tax rate is determined by a competitive auction. Each player submits a sealed bid. The bids are opened, and the higher one becomes the new tax rate. The player making this bid makes the first play at the new tax rate. Then, at each turn, a player may elect to either

- pay the current tax (to his opponent) and play a move on the board, or
- pass at the current tax rate.

After two consecutive passes at the current tax rate, a new auction is held to restart the game at a new and lower tax rate. Legal bids must be strictly less than the prior tax rate. There is also a minimum legal bid, t_{\min} . A player unwilling to submit a bid $\geq t_{\min}$ may pass the auction. The game ceases when both players pass the auction.

If the players submit equal bids, the referee may let an arbitrator decide who moves next. The arbitrator is disreputable: Each player suspects that the arbitrator may be working in collusion with his opponent. Alternatively, when there are equal bids for the tax rate, the referee may require that the next move be made by the player who did not make the prior move. In this way, the referee, at his sole discretion, can implement a “preference for alternating moves”.

In some traditional games, including Japanese Go, the minimum tax rate t_{\min} is defined as 0. When the game stops, the score is tallied according to the position of the board, and the corresponding final payment of “score” is made. This “score” payment is added to the tax payments that the players have paid each other, and to the payment(s) to buy the more desirable color(s) initially. Since all payments are zero-sum, one player’s net gain is his opponent’s net loss. The player who realizes an overall net gain is the winner. In Economic games, a tied outcome occurs whenever each player’s total net cumulative payment to his opponent is zero.

For *Winning Ways*-style games, which include mathematical play on “numbers” beyond the conventional mathematical “stopping positions”, t_{\min} is defined to be -1 . The minimum tax rate can also be defined as -1 for some non-Japanese versions of Go rules.

When such a game ceases, no score remains on the board: the outcome depends entirely on the net cumulative total of payments made between the players.


1. A Demonstration Game

We now hold a demonstration using the Economist's Rules. The players are Yonghoan Kim and David Wolfe.

The referee, EB, gives each player \$5 initial capital, and they get change at the bank, in cash and chips valued as follows:

- \$5, \$1 = US notes
- Q = US quarter
- G = green chip, worth 1/4 of a unit, now set at \$1
- Y = yellow chip, worth 1/16 of a unit

The game they will play is the sum of a Go region and a Domineering region. The starting position of the Domineering region is the empty 4 × 9 board. The starting position of the Go region is shown on the right.



We begin by holding three concurrent auctions. Each player submits one bid for the Vertical dominoes, another for the Black Go stones, and a third for the first move on the sum of both games. All sealed bids are opened concurrently. Here are the results of our demonstration auction:

	YK	average	DW	result
Go Black	$-2\frac{1}{4}$	-2	$-1\frac{3}{4}$	DW plays Black (Left)
Domineering Vertical	$1\frac{1}{2}$	1	$\frac{1}{2}$	DW plays Horizontal (Right)
First Move	$1\frac{1}{4}$	$1\frac{1}{2}$	$1\frac{3}{4}$	DW moves first

For consistency, we decide to negate the Go board so that sides played by YK and DW in each game match our conventional uses of Left and Right, respectively. This is achieved by reversing all colors in the Go position. Here are the auction results restated in terms of -Go:

	YK	average	DW	result
-Go Black	$2\frac{1}{4}$	2	$1\frac{3}{4}$	YK plays Left
Domineering Vertical	$1\frac{1}{2}$	1	$\frac{1}{2}$	YK plays Left
First Move	$1\frac{1}{4}$	$1\frac{1}{2}$	$1\frac{3}{4}$	DW moves first

The following payments are then made:

- YK pays DW \$2 in cash to play Left (Black) on -Go
- YK pays DW \$1 in cash to play Left (Vertical) on Domineering
- DW pays YK \$1.50 in order to make the first move.¹

¹In order to pay the tax for the first move, DW first gets change from YK: \$1 for 1Q + 2G + 4Y. He then puts \$1.50 worth of chips (4G + 8Y) onto the table for YK.

1.1. Erdős' Book. The mathematician Paul Erdős frequently mentions “The Book”, a supernatural reference document that contains all mathematical truths in their most revealing forms. To expedite the remainder of this demonstration, we allow the participants and ourselves to peek into The Book. Figure 1 expresses




Figure 1. Book starting values.

(albeit in perhaps unfamiliar symbols) the Book values for the two starting positions. The values also contain more information than we need to play optimally according to the Economist’s Rules. In addition to a single good move for each player at the current tax rate, for each position we need only its fair market value, μ , and one of its fair tax rates, t (also called the “temperature”). These values are shown on the bottom line of Figure 1.

The remainder of our demonstration will illustrate how players able to peek into The Book can play optimally by following a few very simple common-sense rules, based on two important facts:

- The value of playing Left in any region is μ .
- Although the “size of a move” in any region is not necessarily unique, it may always safely be taken as the temperature of that region.

The *Book Strategy* is as follows:

0. Initially, bid μ for the privilege of playing Left.
1. At every stage of play, define the global temperature of the game as the maximum of the temperatures of its regions. If it is your turn, proceed as follows:
 - (a) If opponent’s prior move has created a region whose temperature exceeds the current tax rate, respond by playing in the same region as the opponent’s prior move.
 - (b) Otherwise, if the temperature of the game is at least as large as the current tax rate, play in a region of maximum temperature.
 - (c) If the tax rate exceeds the maximum regional temperature, pass. If you have any canonical move(s), bid the maximum regional temperature.

The strategy is designed to answer directly this fundamental question:

Should I respond locally to my opponent's prior move?

Restated in Go jargon, that same question becomes

Does my opponent's prior move have *sente*?

Since the Book Strategy focuses on this question, it is also called *Sente*strat. After the demonstration, we will discuss how Sentestrat compares with two other strategies, Hotstrat and Thermostrat.

Once Yonghoan Kim and David Wolfe obtain access to The Book, each of them elects to follow the Book Strategy.




Figure 2. The demonstration game: moves 1 and 2 played according to The Book.

DW begins playing a White Go stone in the center of the bottom row. This is shown as move 1 in Figure 2. The cost of this move was \$1.50, as determined by the initial auction. According to The Book, that bid was a slight error; the true Book value of the initial temperature was $1\frac{15}{32}$. However, White's move 1 has lowered the temperature of the Go region to $1\frac{1}{16}$. The Domineering region has temperature $1\frac{1}{8}$, which is larger. So, now that they are able to reference The Book, after move 1, both players decline to play at a tax rate of \$1.50. Instead, they both bid \$1.125.

Since the bids are tied, the play alternates, and YK is selected to play move 2 at a cost of \$1.125. Since this is less than the \$1.50 worth of chips he has on the table, YK pockets \$0.375 of these chips, and moves the rest of the chips back to DW in payment of the tax on move 2. YK plays move 2 as the Vertical domino shown. According to The Book, this \$1.125 move has increased μ for the Domineering region from $\frac{3}{4}$ to $1\frac{7}{8}$, but it has not changed the temperature,




Figure 3. Book play, moves 3-5.

which remains $1\frac{1}{8}$. So DW is willing to play at the current tax rate of \$1.125, even though that requires him to give all of the chips on the table back to YK.

DW plays move 3 as the Horizontal domino shown in Figure 3. According to The Book, this move has changed μ for the Domineering region from $1\frac{7}{8}$ to $\frac{11}{16}$, a net decrease of $1\frac{3}{16} = \$1.1875$. As this exceeds the \$1.125 that Right paid for the move, Right may appear to have gained \$0.0625 at move 3. However, this apparent gain is ephemeral, because the move also increased the temperature to $1\frac{3}{16}$, which exceeds the current tax rate of $1\frac{1}{8}$. Therefore, following move 3, YK is eager to play move 4, because it is now possible for him to play so as to gain $1\frac{3}{16}$ at a cost of only $1\frac{1}{8}$. The tax on move 4 returns all of the chips on the table to DW, changes μ to $1\frac{7}{8}$, and lowers the temperature back to $1\frac{1}{8}$. DW then returns all of the chips on the table back to YK, and plays move 5. This lowers μ back to $\frac{3}{4}$.

The careful observer will notice that the four-move sequence 2-5 had no net effect on μ . And, although \$1.125 worth of chips moved back and forth across the table several times, after move 5 they all ended up back where they were after move 1.

After move 5, The Book reveals that the Domineering temperature is 1, which is now less than the Go temperature of $1\frac{1}{16}$. So both players then elect to pass




Figure 4. Book play, moves 6–8.

at the tax rate of $1\frac{1}{8}$. The referee holds a new auction to determine who will move next, and both players bid $1\frac{1}{16}$. Because his opponent made the last move at the old tax rate, YK pockets chips worth $\frac{1}{16}$, leaving chips worth $1\frac{1}{16}$ on his side of the table. Since the auction was tied, the referee adopts the alternating move convention and assigns the next move to YK, who pays the new tax rate by moving all of the \$1.0625 worth of chips on the table to DW as he plays move 6 on the Go position, as shown in Figure 4. According to The Book, move 6 increases the local μ to $\frac{9}{16} + 1\frac{1}{16} = 1\frac{5}{8}$, and decreases the Go temperature to $\frac{7}{8}$. The Domineering region is now the hottest game, with temperature $t = 1$. As the old tax rate was $1\frac{1}{16}$, both players elect to pass and bid $t = 1$. DW pockets chips worth $\frac{1}{16}$, because his opponent made the last move at the old tax rate. The value of the chips on the table is then \$1, on DW's side. The referee assigns the next move to DW, who pays the chips back to YK and plays move 7, followed by YK's move 8 and DW's move 9 (in Figure 5). The net effect of the three-move sequence 7–9 was a payment of \$1 of chips from DW to YK, in return for a \$1 change in the value of the Domineering μ in DW's favor.

After move 9, both players pass and bid $\frac{7}{8}$. The tax rate decreases by $\frac{1}{8}$, which YK pockets, leaving \$0.875 worth of chips on the table. These remaining chips




Figure 5. Book play, moves 9–12.

are transferred to DW as YK plays move 10. This move increases the Go value of μ by the same amount as the tax paid, $2\frac{1}{2} - 1\frac{5}{8} = \frac{7}{8}$.

After move 10, both players pass and bid $\frac{3}{4}$. The tax rate decreases by $\frac{1}{8}$, which DW pockets, leaving \$0.75 worth of chips on the table. These remaining chips are transferred to YK as DW plays move 11 on the Domineering board. This move decreases the Domineering value of μ by the same amount as the tax paid, $-1 - (-\frac{1}{4}) = -\frac{3}{4}$.

After move 11, both players pass and bid $\frac{1}{2}$. The tax rate decreases by $\frac{1}{4}$, which YK pockets, leaving \$0.50 worth of chips on the table. These remaining chips are transferred to DW as YK plays move 12 on the Go board. This move has the effect of increasing the Go value of μ by the same amount as the tax paid, $3 - 2\frac{1}{2} = \frac{1}{2}$.

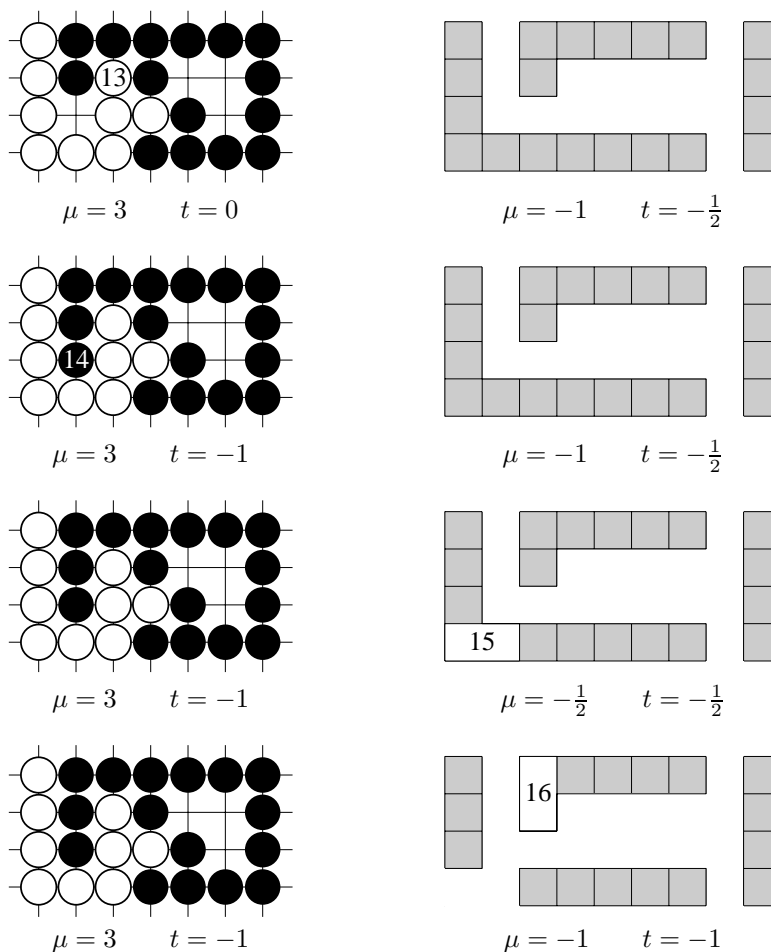


Figure 6. Book play, moves 13–16.

According to The Book, after move 12, the temperature of the Go region is 0 and the temperature of the Domineering region is $-\frac{1}{2}$. So both players pass at the tax rate of $\frac{1}{2}$, and submit bids of 0 for the new tax rate. DW pockets all \$0.50 worth of chips on the table. The new tax rate is now 0, which equals the value of the chips on the table.

DW plays move 13 in Figure 6, which is the first move at the tax rate of 0. YK responds with move 14, at the same tax rate of 0.

After move 14, The Book reveals that both the Go position and the Domineering position have negative temperatures. So both players would prefer not to play. In fact, neither will play unless his opponent pays him to move! So each player bids $-\$0.50$. This means that neither player will move unless his opponent pays him \$0.50. The new tax rate is negative! So DW *collects* \$0.50

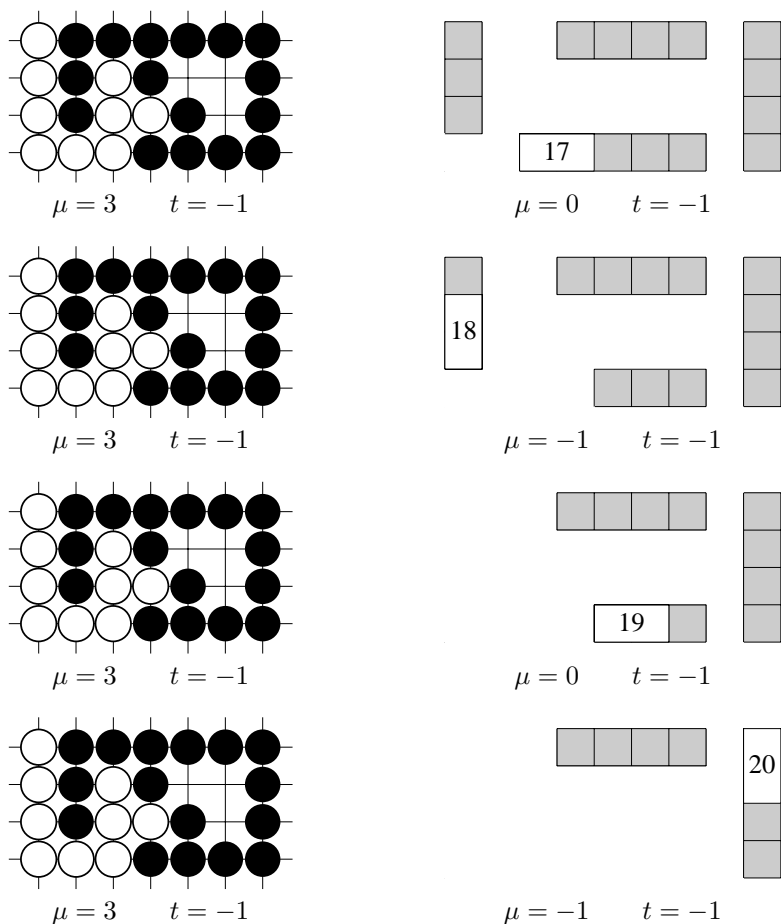


Figure 7. Book play on more integers, moves 17–20.

in cash from YK as he plays move 15. But YK responds with move 16, in return for which he gets his \$0.50 back.

After move 16, The Book reveals that both positions have temperature equal to negative one. Both players pass at the tax rate of $-\frac{1}{2}$, and submit bids of -1 to set the new tax rate.

DW collects \$1 cash from YK for move 17 in Figure 7, but YK collects it back for move 18. The players continue to “cash in” their positions, at the rate of \$1/move, on moves 19 and 20 in Figure 7, then 21, 22, and 23 in Figure 8.

According to Japanese-style Go rules, the play ceases after move 23. The Go position is scored as +3 points for Black, and White is required to play Black \$3 to settle this score.

According to the mathematized (Universalist) Go rules, the game continues after move 23. White submits no bid. Black submits the minimum bid that the

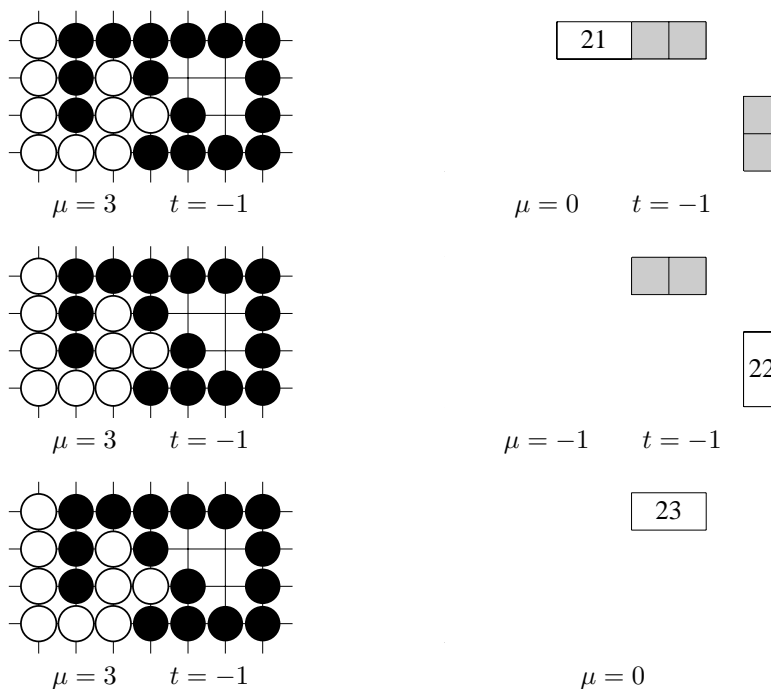


Figure 8. Book play on more integers, moves 21–23.

rules allow, which is -1 . So Black wins the right to play at a tax rate of -1 . In Figure 9, he plays move 24 as a Black Go stone in the middle of the three empty spaces, and collects \$1. (According to Universalist Go rules, the Black Go stones are now “immortal”, and cannot be captured even if Black fills in both of the visible eyes.)

Again, White passes and submits no bid. So Black again submits the minimum allowable bid of -1 , and plays move 25 as another Black Go stone, for which he again collects \$1 from White. Again White has no option but to pass and submit no bid. Black again submits the minimum allowable bid of -1 , wins the auction, and plays move 26 as another Black Go stone, for which he collects his final \$1 from White.




Figure 9. Book play on integers remaining after move 23.

After move 26, both players pass. Neither player has any remaining legal move, so neither can bid. The game finally ends, and our demonstration is concluded.

1.2. Theorem: Book Play Ensures No Monetary Loss. After the initial auction has determined who plays Left and who plays Right, then either player may elect to use the Book Strategy to select his subsequent tax rate bids and plays. The following theorem asserts that this strategy is optimal for games played according to the Economist's Rules:

THEOREM. *If Left plays Book Strategy, she ensures a total net gain of at least μ dollars. If Right plays Book Strategy, he ensures that Left's total net gain will be at most μ dollars.*

2. Thermography

Thermography may be viewed as a methodology for deriving the Book values of μ and t . Classical thermography for loopfree games is presented in [Conway 1976] and *Winning Ways*. In this paper, we present an extended version of this methodology, which also handles loopy Go positions called kos. To this end, we must first extend the Economist's Rules to deal with kos.

2.1. The Economist's Ko Rules. The game of Go includes some loopy positions. By far the most common such positions are 2-cycle loops. Locally, White captures a stone, Black captures it back, then White again captures Black's stone, etc. These positions are called kos. All popular dialects of Go rules prohibit such loopiness, by banning the immediate recapture of a single stone if it would return the board to the same position it reached after mover's prior turn.

The extension of thermography to cover kos is facilitated by the following extensions of the Economist's Rules:

For each region of the game, a specified player is designated as the local *Komaster*. In principle, the privilege of being Komaster might be determined by local auctions, held before the play begins. If the global game contains several regions with potential kos, then it is quite possible for Left to be Komaster of some regions, and Right to be Komaster of others. (These rules do not attempt to model positions such as the classical triple ko, in which multiple interdependent kos appear within a single region.)

According to the Economist's Rules, kos are resolved by overruling the preference for alternating moves. Any move that repeats an overall game position is allowed. If the mover is not the Komaster of the region wherein the move is made, there are no further restrictions. However, if the local Komaster makes a move that repeats the global board position (and, necessarily, the local regional position), I call this local ko position *critical*. The Komaster, who just moved to the critical ko position, is compelled to make another local move immediately.

This rule for resolving *ko* overrides the referee's normal preference for alternating moves. Instead, the Komaster makes two moves consecutively, to and from the critical *ko* position, and the Komaster pays two taxes (at the same tax rate), one tax for each move.

2.2. Bottom-Up Thermography. As mentioned earlier, thermography may be viewed as a methodology for calculating the book values, μ and t . The version we now present handles *kos*. It also gives the classical results in the special case of loopfree games.

Trajectories. We call the basic data structure of thermography a *trajectory*. A trajectory is a piecewise linear real function of a real variable. Its input argument, t , is called the *temperature*, or the *tax rate*, and can range from 0 to $+\infty$ (or from -1 to $+\infty$, according to the extension made in Section 4.2.) The trajectory's output value is a real number, called v or μ . Within each linear piece of the trajectory, the value of the derivative dv/dt is constant and equal to an integer. In elementary calculus, it is conventional to plot the independent variable on the horizontal axis and the dependent variable on the vertical axis. However, in thermography it is conventional to rotate the conventional calculus plots counterclockwise by 90° , in conformity with a long-standing tradition introduced by Conway. The temperature, t , is plotted along the vertical axis, with increasing values of t corresponding to higher positions along the axis. The value, v , is plotted along the horizontal axis, with the convention that increasing positive values are plotted to the Left; large negative values, to the Right. This facilitates a more direct correspondence with the formal expressions for most hot games, in which Left moves toward more positive values and Right moves toward more negative values.

A thermograph is a pair of trajectories, called stops or scores or *walls*. In this exposition, we use the term "wall". The Left wall, LW, is always at least as great as the Right wall, RW.

Scaffolds, Hills, and Caves. The final stages of the construction of a thermograph begin with another pair of trajectories LS and RS, called *scaffolds* (Figure 10, left). In the most general case, the Left and Right scaffolds might cross several times. A temperature interval in which the Left scaffold exceeds the Right scaffold is a *hill*, or *solid region*. A temperature interval in which the Right scaffold exceeds the Left scaffold is a *cave*, or *gaseous region*. Within a hill, the Left wall equals the Left scaffold and the Right wall equals the Right scaffold. However, within a cave, the Left and Right walls coincide in a local trajectory called the *mast*.

The lowest region of a nondegenerate Thermograph is necessarily a hill. The interval along the v -axis, or "ground", between the walls of this hill is the *base* of the hill. The point at which the two scaffolds cross at the top of the hill is both the *summit* of the hill and the *drain* of the cave above it. If the scaffolds

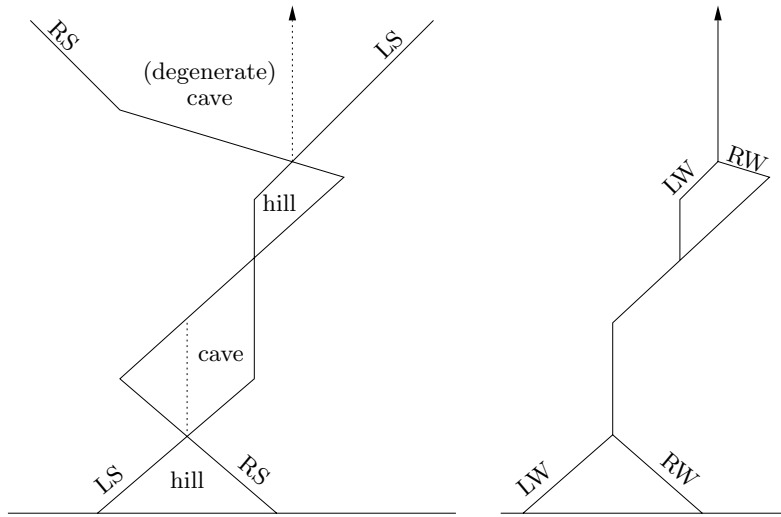


Figure 10. Scaffolds (left), and the thermograph constructed therefrom (right).

cross again at the top of the cave, that point is called the *chimney* of the cave, and the *fulcrum* of the hill above it.

Normal Masts. Within any cave, the normal mast can be easily constructed by following the path of an ideal balloon from the drain to the chimney. When free to do so, the balloon rises vertically. If it hits a scaffold, the balloon follows that roof scaffold upward until it is either again free to rise vertically or until it reaches the chimney, at which point the mast ends and the Left and Right walls again become different, starting at the fulcrum of the adjacent hill. This is shown on the right in Figure 10. Moving upward from $t = 0$, this thermograph has a hill, a cave, another hill, and then an infinite gaseous region (a degenerate cave) through which the mast rises vertically to infinity. All hot games have thermographs with a hill at $t = 0$ and a vertical mast as t approaches infinity.

The walls give the values of the economic game with a current tax rate of t .

If the thermograph is a hill region at temperature t , then each player desires to move first. By doing so, he can ensure a value given by the corresponding wall. If the thermograph is a cave region at temperature t , then at least one player prefers to pass until the tax rate becomes lower. If the mast is vertical, then both players prefer to pass, and neither will play until the tax rate is lowered. If the mast coincides with a roof of the cave which is a Left scaffold, then Left prefers to play and Right prefers to pass. If it is Right's turn, he will pass and propose a tax cut, but Left will then reject Right's proposal. Left will instead exercise his right to move at the old, higher tax rate. In a subsequent subsection we will see a specific example of this surprising phenomenon, where it is to a player's advantage to insist on playing at a higher tax rate rather than at a lower one.

DEFINITIONS. We define the *initial temperature of a game* G , as the temperature of the base of the highest mast of its thermograph. A game G is said to be *active at temperature* t unless its thermograph at t is a vertical mast. In that case, we say that G is *dormant at* t . If G is dormant at t , then the temperature of the base of this mast is called the *activation temperature of* G at t .

Pass-Banned Masts. The constructions of thermographs must reflect the Economic Rules that resolve kofights. As described in Section 2.1, these rules do not allow the Komaster to pass from a critical position. Under such circumstances, the masts within each cave degenerate into the maximum (Right) scaffold if Right is not allowed to pass, or the minimum (Left) scaffold if Left is not allowed to pass.

Recursion for Thermographs of Games with No Immediate Ko. Let G be a game represented in canonical form, $G = \{G^L \mid G^R\}$. where G^L and G^R are sets of games whose thermographs are already known. We first consider the case in which G itself is not part of any loop, even though some of the positions in G^L and G^R may contain kos.

Under these circumstances, the scaffolds of G are defined as

$$\begin{aligned} \text{Left scaffold of } G_t &= -t + \max_{G^L}(\text{Right wall of } G_t^L), \\ \text{Right scaffold of } G_t &= t + \min_{G^R}(\text{Left wall of } G_t^R) \end{aligned}$$

Thermographs of Kos. Figure 11 shows a simple 33-point ko and the graph of its legal moves.




Figure 11. Simple 33-point Ko.

Let G and H be a pair of games that represent the two states of a ko. G is a Left follower of H , and H is a Right follower of G . In the general case, G and H may also have other followers; thus we may have

$$G = \{V \mid W, H\}, \quad H = \{G, X \mid Y\}.$$

For such a ko, we construct two pairs of thermographs. \hat{G} and \hat{H} assume that Left is the regional Komaster; \check{G} and \check{H} assume that Right is the regional Komaster.

From \hat{H} , Left may play to a close relative of G , which we call \hat{G} .

The relevant pair of thermographs are computed by working through the following three equations, from the bottom up:

$$\begin{aligned}\hat{G} &= \{V|W, \hat{H}\}, \\ \hat{H} &= \{\hat{G}, X|Y\}, \\ \hat{G} &= \{V|W\}, \text{ with proviso that Left cannot pass immediately.}\end{aligned}$$

The motivation behind these equations is as follows: Since Left, as komaster, can win the ko, she may effectively prevent Right's move from \hat{G} back to \hat{H} . However, the price of this prevention is the constraint that if Right elects to pass because of the koban, Left is forbidden the privilege of passing immediately from \hat{G} . Such a pass would permit Right to recapture the ko at a lower tax than Left paid to get there, and to return to the same board position as \hat{H} . However, this new position would be accompanied by a ko state less favorable to Left than \hat{H} . So if Left is competent, we should never see the sequence of play in which Left moves from \hat{H} to \hat{G} , and later passes when position \hat{G} is still available. If Left intends to do that, she could do at least as well by passing immediately from \hat{H} .

Similarly, to find the thermographs of \check{G} and \check{H} , we may use the next three equations, again from the bottom up:


$$\begin{aligned}\check{H} &= \{\check{G}, X|Y\}, \\ \check{G} &= \{V|W, \check{H}\}, \\ \check{H} &= \{X|Y\}, \text{ with proviso that Right cannot pass immediately.}\end{aligned}$$

A similar methodology enables us to compute thermographs for iterated kos. For example, suppose thermographs A and E are known, and that

$$B = A|C, \quad C = B|D, \quad D = C|E.$$

Then, if we are interested in the games in which Left wins these kos, we may compute thermographs of \hat{B} , \hat{C} , \hat{D} , \hat{C} , and \hat{B} , in that order.


Here are the thermographs derived by this methodology for the positions of Figure 11, when White is komaster. (Naturally, $\check{A} = A$, with Right as komaster; $\check{B} = B$, with Right as komaster; $\check{B} = \check{B}$ with Right not allowed to pass.)



A^L is a vertical mast at $v = 24$. B^R is a vertical mast at $v = -9$. The pass-banned \check{B} is a mast of slope +1 starting from $t = 0$, $v = -9$. The Left scaffold

of \check{A} is obtained by subtracting t from A^L . It is a ray with slope -1 , from $t = 0, v = 24$. The Right scaffold of \check{A} is obtained by adding t to $\check{A}^R = \check{B}$. It is a ray with slope $+2$ from $t = 0, v = -9$. The two scaffolds intersect at $t = 11, v = 13$. For $t \leq 11, LW(\check{A}) = LS(\check{A})$ and $RW(\check{A}) = RS(\check{A})$. For $t \geq 11, LW(\check{A}) = RW(\check{A}) = 13$.

The thermograph for \check{B} of Figure 11 is constructed as follows:



$RS(\check{B})$ is a ray of slope -1 from $t = 0, v = -9$. $LS(\check{B})$ is obtained by subtracting t from $RW(\check{B}^L) = RW(\check{A})$. For $t \geq 11, LS(\check{B}) = RS(\check{B})$, but $LS(\check{B})$ changes direction at $t = 11$. Above $t \geq 11, LS(\check{B})$ and $RS(\check{B})$ form a cave, in which $LW(\check{B}) = RW(\check{B}) = 2$. For $0 \leq t \leq 11, LW(\check{B}) = RW(\check{B}) = -9 + t$.

The derivation of thermographs for \hat{A} and \hat{B} is directly analogous. Here are the results:




Figure 12, left, shows a Go position called a *seki*. This position is terminal because competent players will pass. Mathematically, its value is

$$24|-10||26|-8 = 0.$$

Its thermograph is a vertical mast at $v = 0$.

Figure 12, middle, shows G, a famous position from traditional Go literature called the “thousand-year Ko”. Its game graph is shown on the right in the same figure. From G, White can play either to the seki of value 0 just mentioned, or to B, the ko of Figure 11. Left has two similar options, whose thermographs are superimposed in Figure 13. From \check{G} , if $t < 5$, Left prefers to play to the seki; but if $t > 5$, Left prefers to start the ko. Figure 14, left, shows the result of




Figure 12. The seki of value 0 (left) is one option for White in the “thousand-year Ko” (middle), whose game graph is shown on the right.




Figure 13. Thermographs of black's options from \check{G} .

translating the maximal Right wall of \check{G}^L into the Left scaffold of \check{G} . Figure 14, right, shows $RS(\check{G})$.

The thermograph of \check{G} , shown in Figure 15, left, illustrates some features of the thousand-year ko that most intermediate Go players find counterintuitive. If the global temperature exceeds 11, both players will prefer to play elsewhere rather than on \check{G} . The activation temperature of \check{G} is 11. At temperatures between 11 and 7, White prefers to play elsewhere, but Black is eager to start the ko from \check{G} . White will win this ko, but doing so costs him two moves and it costs Black only one move to start the ko. Thus, at temperatures between 11 and 7, the higher the global temperature at which Black can start the kofight, the happier she is.




Figure 14. Left and Right scaffolds of \check{G} .



Figure 15. Left: Thermograph of \check{G} . Right: Thermograph of \hat{G} .

At global temperatures between 3 and 7, \check{G} is again inactive. Both players prefer to play elsewhere.

At global temperatures below 3, both players are eager to play on \check{G} . White would like to start and win the ko, and Black is eager to avoid this loss by playing to the seki.

The thousand-year ko has two activation temperatures, and a region wherein one player prefers to pay a higher tax rate rather than a lower one. These features are unusual, and dependent on the assumption that White is Komaster. The thermograph of \hat{G} , shown on the right in Figure 15, is much more common-looking. At high temperatures, the value of being Komaster on G is $7\frac{2}{3} - 1 = 6\frac{2}{3}$ points. This contrasts sharply with the common ko of Figure 11, where being komaster has zero value at high temperatures.

2.3. Overview of Thermography.

Combinations and Sums. Thermographs support combinations. The thermograph of A contains just enough information so that, when combined with the thermograph of B or C , one can obtain the thermographs of $A|B$ and of $C|A$.

Thermographs do *not* support addition, as is evident from the following example. Let $A = 1|-1$, $B = \{1, 1 + A|-1\}$, $C = \{1, 1 + A|-1, -1 - A\}$.

Then A , B , $-B$, and C each has the thermograph shown on the left in Figure 16, which we call X . Yet $A + B$ has thermograph Y ; $A - B$ has Z ; $A + C$ has X ; and $A + A$ has 0 .




Figure 16. Thermographs don't behave well under addition.

For a general loopfree game, the shape of the thermograph at the base of its mast depends on the sign of the infinitesimal by which the game differs from

a number at its freezing point, and the information contained in the values of these infinitesimals is not present in the thermographs. We will return to this topic in Section 2.5.

However, at sufficiently high temperatures, all thermographs become masts, and their mast values do add. Therefore, knowing the thermographs of all of the summands *is* sufficient to determine the thermograph of the sum *at sufficiently high temperatures*.

Means and Temperatures. The μ -value of the initial, infinite, vertical mast, which we call $\mu(G)$, is the “count”, or the “mean value” of G .

The temperature at the base of this mast, which we call $t(G)$, is the “size of the move”, or the “temperature” of G .

Placid Kos and Hyperactive Kos. A game G is said to be “hyperactive” if $\mu(\hat{G}) > \mu(\check{G})$. Such a game G necessarily contains a position K that is a ko. It is not the size of K itself that makes G hyperactive, but special circumstances which include the fact that K is a follower of some position of G that is cooler than K .

A game G is said to be “placid” if $\mu(\hat{G}) = \mu(\check{G})$. If G is placid, then when played according to Economic Rules, $\mu(G)$ does not depend on kothreats. Many kos, even many rather big kos, turn out to be placid.

Stable and Unstable Positions. Thermographs with more than a single hill and a single (degenerate) cave are relatively rare. In the common case, stable and unstable positions can be defined relatively simply, as follows:

DEFINITIONS. A position H of a game G is *stable* in G if all ancestors of H in G have temperatures greater than H 's. H is *unstable* or *transient* in G if it has an ancestor of lower temperature. If H is not unstable, but has one or more ancestors of temperature equal to its own, it is called *semistable*.

In the more general case, in which some positions may have thermographs with multiple hills and caves, a single position may have multiple activation temperatures, and so a more refined definition is required:

REFINED DEFINITIONS. G is *stable in* G , and its activation temperature is the base of its infinite mast. H is *stable in* G if its thermograph is a vertical mast at the activation temperature of its nearest stable ancestor.

Semi-stable positions can be treated as stable or as unstable. Treating them as stable yields the simplest and most convenient economic analyses. Treating them as unstable leads to more refined, more complicated algorithms that exploit the calculus of infinitesimal games to get the last move at each t , whenever possible.

Sentestrat players will choose to play “in another region” “elsewhere” only when the region just played in has reached a stable position.

Sente and Gote. From the perspective of Sente strat, an opponent's move that increases the temperature has sente; it requires our immediate response. However, this move and its response might begin a longer sequence of moves through several transient positions before the game reaches the next stable position. The question of whether the overall sequence had sente or gote depends on whether the length of this alternating sequence of moves was odd or even, and that question can usually be answered by examining the thermograph of the originating stable state. Generally speaking, if, just below the base of its relevant mast, both walls branch outward, then the initial move by either player is *gote*. If Left's wall drops vertically but Right's branches outward, then Left's initial move is *sente* but Right's is *reverse sente*. If Right's wall drops vertically but Left's branches outward, then Right's initial move is *sente* and Left's move is *reverse sente*.

These correspondences between the Go player's view of sente and gote and the thermograph are correct whenever the relevant positions have significantly different temperatures. However, when there are intermediate semi-stable positions, the distinction between sente and gote may become blurred.


Mainlines and Sidelines. Let n be a large integer, and suppose that two gurus play the sum of n copies of a game G . For each H that is a position of G , let $\#(H)$ be the number of copies of G that pass through the position H . If $\#(H)$ approaches infinity as n approaches infinity, then H is a "mainline" position of G . But if $\#(H)$ remains bounded as n approaches infinity, then H is a "sideline" position of G .

If G is a traditional gote position, then $\#(G^L) = \#(G^R) = n/2$, and both G^L and G^R are mainline positions. However, if G is a position in which Left's move is sente and Right's move is reverse sente, then, depending on who plays first from $n.G$, $\#(G^R) = 0$ or 1 , and $\#(G^L) = n$ or $n - 1$. In either case, G^R is a sideline position of G .

Thermographs of sideline positions have no effect on the means of any of their ancestors, although they often do affect the temperature(s) of some of their near ancestors.

We illustrate these notions with an example.

Let A be the starting position on the right, before any of the numbered stones have been played. Then moves 1, 2, 3, and 4 show a main line of play, through positions we'll call B , C , D , and E . Position E is the terminal position after all four numbered stones have been played.



Position F is the position if White rather than Black had played the first stone at 1. The means and temperatures of all dominant positions are shown in Table 1.

The only stable mainline positions are A , C , and E , each having mean 1. The mean of A is unaffected by any small change in the mean or temperature of any other position. Hence, to verify that the mean of A is 1, it is unnecessary to

Position	μ	t	Stable?	Mainline?
A	1	3	yes	yes
B = A ^L	8	7	no	yes
B ^L	15	-1	yes	no
C = B ^R	1	1	yes	yes
C ^L	2	-1	yes	no
D = C ^R	$-\frac{1}{2}$	$1\frac{1}{2}$	no	yes
D ^R	-2	0	yes	no
E = D ^L	1	-1	yes	yes
F = A ^R	-2	1	yes	no
F ^L	-1	-1	yes	no
F ^R	$-3\frac{1}{2}$	$1\frac{1}{2}$	no	no
F ^{RR}	-5	0	yes	no
F ^{RL}	-2	-1	yes	no

Table 1. Means and temperatures of dominant positions.

know the values of other means and temperatures precisely; appropriate bounds are sufficient.

In general, a stable mainline gote position will have both a mainline Black follower and a mainline White follower. On the other hand, stable mainline sente positions have only one mainline follower. In this example, the line of play shown is the unique main line. Notice that the mainline moves do not alternate colors; White plays both moves 2 and 3.

If we played a sum of n copies of A concurrently on n different Go boards, then the local line of play would follow the main line shown above on all but one or two boards. If White gets the first move overall, he plays one exceptional board from A to a sideline position, F. But Black then plays another board from A to B, and White responds by continuing to C. That happens on all $n - 1$ boards, successively. Black then might play a second exceptional board from C to C^L. But White then plays another C to D, and White responds by continuing to E. That sequence, from C to D to E, then happens on all boards, successively, until all positions are terminal.

The sum of n copies of A is only one example of a rich environment, which we shall explore further in Section 4.3. If a single copy of A is played in any sufficiently rich environment, the play within A will follow the main line shown above. Heuristically, the reason that Black rather than White will make the first move from A is that Black can move there at any tax rate between 7 and 3; White cannot afford to play A until the tax rate is 3. After A has been played to B and C, both players make several moves elsewhere if possible. Eventually, any time after the tax rate drops below 1.5, White can play at C. Since Black cannot afford to play there until the tax rate is 1, White will get in his move

first, unless the background is so impoverished that there are no suitable moves “elsewhere”.

Smaller Thermographic Databases of G. Given a game, G , and its purported thermograph, how much information about the positions of G is needed to verify that the purported thermograph of G is correct? Evidently, considerably less than the thermographs of all of the positions of G . Many positions of G are transient. Often the only information needed about many of the followers of transient positions is that they are “sufficiently hot”, and this bound can easily be quantified. Precise thermographs of many sideline positions are also unnecessary; it is required only that portions of such thermographs fall between specified bounds.

How much information is needed to verify the purported mean of G ? Evidently, even less than is needed to verify G 's thermograph.

2.4. Top-Down Thermography. “Top-down” thermography promises to be much more efficient than “bottom-up” thermography. The latter computes the details of many walls of positions that the former could ignore.

A proposed top-down approach to finding the thermograph of a position G starts by considering a reasonably large set of lines of play from G , such as $G^{\text{LRL}_1\text{LRLRRLR}}$. Most relevant lines do *not* follow alternating play. Indeed, empirical evidence indicates that Go positions, such as the thousand-year ko of Figure 12, which have different dominant Left options depending on the temperature, are fairly rare. So the main reason that there are typically many relevant lines of play from G is that there may be many relevant sequences of Left and Right plays, rather than choices between H^{L_1} and H^{L_2} .

The proposed approach to computing the “top-down” thermograph of G assumes that the given set of lines of play contains the “complete” set of G 's stable positions. This assumption allows a program to infer that missing positions are sufficiently hot.

One is often interested in only the high-temperature portions of thermographs. For this reason, it would eventually be desirable to have programs that do *all* thermographic computations from a top-down perspective. Each trajectory might be stored as a list of slopes and breakpoints, listed in top-down order. A better programming strategy is to view each trajectory as a *function* whose outputs are a list of slopes and breakpoints, generated in the top-down order. This function is able to call other functions and itself, recursively, as needed, in order to pursue the top-down computation of thermographs. The programming goal is to get the next output as quickly as possible, by looking no further than necessary.

Thermography Justifies Some Go Maxims. Many of the standard Go maxims are equivalent to some of the common special cases that a top-down thermography program will encounter. For example, if we let $\tau = \frac{1}{2}(\mu(G^{\text{L}}) - \mu(G^{\text{R}}))$, and




Figure 17. Thermographs of loop-free infinitesimals.




Figure 18. Thermographs of numbers.

if $t(G^L) \leq \tau$ and $t(G^R) \leq \tau$, then G might be said to be “gote”, and we have $\mu(G) = \frac{1}{2}(\mu(G^L) + \mu(G^R))$, and $t(G) = \tau$. This condition is very common.

One type of situation in which thermography yields much more precision than traditional Go folklore is the “double sente” position, in which $t(G^L)$ and $t(G^R)$ are both substantially larger than $LW(G^L) - RW(G^R)$. Traditional Go maxims indicate only that a move from such a position G is very desirable, presumably relative to the naive approximation of movesize as $LW(G^L) - RW(G^R)$. This is true if the overall temperature is low and the position has just arisen.

2.5. Subzero Thermography. Although thermographs are conventionally plotted for temperatures ranging from 0 to infinity, they can also be plotted for negative temperatures. Temperatures infinitesimally less than zero appeared in *Winning Ways*, page 151, as “Foundations of Thermographs”. Since the Economist’s Rules allow temperatures as low as -1 , thermographs can be extended into that same region. The four possible thermographs for an infinitesimal game are shown in Figure 17.

The astute reader will notice the correspondence between Figures 16 and 17.

At negative temperatures, the thermographs of noninteger numbers also become active, as illustrated in Figure 18.

The half-ish thermographs are shown in Figure 19.

If v is a number, at positive temperature, the thermograph of any v -ish game is a vertical line emating upward from the value v , which can be any rational number whose denominator is a power of 2. These values form a dense set on the line $t = 0$. However, the possible values of thermographs become sparse at negative temperatures. Figure 20 shows the superposition of the thermographs




Figure 19. Thermographs of games infinitesimally close to $\frac{1}{2}$.




Figure 20. The foundations of number-ish thermographs.

of all number-ish loopfree games, in the temperature range between 0 and -1 . Figure 21 illustrates where the thermographs of two particular numbers appear in Figure 20.

If one views an arbitrary conventional thermograph from a positive temperature perspective, it may appear to be resting on “solid ground”, consisting of the v -axis at $t = 0$. However, if one probes deeper as in Figure 20, one finds that at $t = -1$ all thermographs rest on a discrete set of pilings located at the integer values. At $t = -\frac{1}{2}$, there are twice as many pilings located at half-integers, etc.

An earlier, less universal, version of Figure 20, based on *overheating* rather than on the Economist’s Rules allowing negative temperature, appeared on page 172 of *Winning Ways*.




Figure 21. Underground thermographs of $\frac{1}{2}$ and $-\frac{3}{8}$.

3. Theorems and Strategies

Thermography always gives the unique value for the count, $\mu(G)$, and a usable value for the size of the move, $t(G)$. We now look at some of the theorems that these quantities satisfy.

3.1. Theorem: Thermography Determines Fair-Market Values for Games Played by Economic Rules. We assert that thermography determines fair economic values for μ and for the initial tax rate t . These values are “safe” to bid. For a single game (no sum), the proof follows recursively from the “Economist’s Rules” and the way the thermograph was constructed.

3.2. Theorem: μ -Values Add; Temperatures Maximize. The mast-value, μ , is analogous to the mean in probability. The temperature, t , is analogous to the standard deviation. As in probability, the mean value of a sum of things is the sum of their mean values. But the dispersion about this mean is much more constrained in combinatorial game theory than it is in probability. In probability, the standard deviation of the sum of n objects typically grows as the square root of n . In combinatorial games played by the Economist’s Rules, the dispersion is zero. In combinatorial games played according to the normal rule of alternating play, the dispersion of the sum about its mean is constrained by an upper bound independent of n . For sums containing no positions with hyperactive kos, this bound is proportional to t , and the constant of proportionality depends on the maximal iterative depth of any of the placid kos.

When multiple hyperactive kos are present, we need to define “independence” of regions. This can be done by constructing each local position on a different board, and then playing all boards concurrently. At each turn, a player selects any board and makes one move on that board. When the game ends, each player receives a single score, equal to the total score on all of the boards. Each ko that Left can win is placed on a board with a plentiful supply of Left kothreats; each ko that Right can win is placed on a board with a plentiful supply of Right kothreats. We further define the ko rule to prohibit any recapture of ko unless there is a change of position *on the same board*, so that no move on another board constitutes a kothreat and different boards are really independent of each other.

3.3. Sentestrat, Thermostrat, and Hotstrat. “Hotstrat” is an intuitively obvious strategy for playing a sum of games. It says “always move on a hottest summand”. Unfortunately, we shall see that in some (uncommon) cases, this strategy turns out to yield losing results. This defective behavior of Hotstrat is evaded by either of two slightly refined strategies, called “Thermostrat” and “Sentestrat”. If the current tax rate is t , and there are several current positions whose thermographs are hills (or summits) at temperature t , then these three strategies choose among those options in slightly different ways:

- *Hotstrat*: Move in a region whose temperature is maximal.
- *Thermostat*: Move in a region whose hill is widest at temperature t .
- *Sentestrat*: If your opponent has just played in one of these hill regions, respond directly by playing in the same region in which he just moved.

All three strategies essentially agree under other circumstances. If all thermographs are caves at temperature t , then pick one of them whose mast touches your own scaffold at temperature t , if any such exist. When you are about to play a move of this type, you should *decline* any tax cut proposal that may be pending. If the only hill move at the current temperature is banned by ko, then *pass* but do not lower the bid tax rate. If there are no moves of any of the previously mentioned types, then you should *pass* and propose a lower tax rate. Sentestrat proposes the tax rate equal to the next lowest value at which you would be willing to move, namely, the drain of a cave or the point at which a vertical mast touches your (necessarily jagged) scaffold inside the cave. Thermostat calculates other trajectories in hopes of getting a stronger result. Thermostat computes the sum of all of your opponent's walls from Sentestrat's proposed tax rate on downwards, and to this sum it adds the width of the widest cave at each temperature. It then picks the temperature at which this trajectory is most favorable. Often, that is the same temperature as Sentestrat's, but on some examples (including one shown in *Winning Ways*), Thermostat takes advantage of an opponent's recent mistake to outperform Sentestrat, even though both do equally well against a good opponent.

Hotstrat's Failure. Although it usually recommends the same moves as Sentestrat and Thermostat, Hotstrat plays poorly in some situations. An example is the following sum $A + B$, as played by Left, going second:

$$A = 1 \parallel 10 \parallel 0 \parallel -20 \parallel -21, \quad B = 1 \parallel 0 \parallel -18.$$

Both A and B are loopfree and have temperature 1. Right plays A to A^R . All three strategies play A^R to $A^{RL} = C = 10 \parallel 0 \parallel -20$. Right then plays B to $B^R = D = 0 \parallel -18$.

Left now faces the crucial decision, whether to play on C or on D . Sentestrat and Thermostat both play D , after which Right plays C to C^R and Left to C^{RL} , giving a total score of 0. Hotstrat, on the other hand, computes the temperatures: $t(C) = 10$, and $t(D) = 9$. Ignoring the fact that both of these temperatures are far above the current tax rate (which has been 1 since the beginning of this example), Hotstrat plays C to 10, after which Right plays D to -18 , giving a total score of -8 . This failure of Hotstrat occurs under either the Economist's Rules, or under the Conventional Rules requiring alternating play without taxes.

4. An Economic Guru Kibitzes a Conventional Game

Sentestrat and Thermostrat both ensure perfect play according to economic rules. Either strategy can also be used in conventional games, in which there are no taxes and alternating play is required. In loopfree games, these strategies ensure that the first player will attain a stopping position at least as good as the mean.

When there are many summands and the temperature is t , a good player can often improve that score by about $\frac{1}{2}t$. To gain more understanding of this situation, it is helpful to view a conventional game through the eyes of a kibitzing economic guru.

4.1. Hard and Soft Currencies. It is convenient to refine the Economist's Rules by introducing dual currencies, cash and chips. The rules that specify which currency must be used for which types of debts are as follows:

Debt	Required Form of Payment
Privilege of playing Left	Cash
Privilege of playing First, at beginning of game	Chips, half of which must be bought from the opponent for cash
Tax on mover during game, while tax rate is positive	Chips
Payment to mover while tax rate is negative	Cash

Thanks to the footnote on page 367, all payments made during our demonstration followed these rules.

Except for kos, conventional play may be viewed as economic play with a preference for alternating moves, and with the proviso that chips are worthless.

In a typical economic game, the pile of chips on the table moves back and forth between the two players frequently. Only when the tax rate is lowered is a player able to move some of these chips off of the table and into his pocket. Economically, this is an apparent gain for the first player to move at the new and lower temperature. The player who made the last move at the old, higher temperature sees an immediate loss of chips. However, if she is following Sentestrat, which ensures no net overall economic loss, then she will eventually receive a compensating payment later. If this occurs after the tax rate has become negative, then that payment will be in cash.

So the conventional player who gets the last move at some temperature would appear to receive a benefit equal to the change of temperature.

There is an alternative, heuristic view that leads to similar conclusions without assuming that the chips are worthless. In this view, the big pile of chips on the table at early stages of the economic game is very likely to be pocketed in many

small decrements. These decrements are likely to be divided sufficiently evenly between the two players that, to a good approximation, the chips can be ignored.

4.2. The Economist's Forecast. In his role as a kibitzer of a conventional game, our economic guru computes its thermograph. He can also keep track of the temperature at each stage of the game. Before the game begins, the economist presumes the prehistoric (big bang?) temperature, $t = +\infty$. Then, at each stage of the game, the kibitzer determines whether the game is active or dormant at the temperature he has determined. If dormant, he lowers his temperature as little as needed to make the game active.

Our economic kibitzer also computes an economic *forecast* based on these observations and computations. If it is Left's turn to play from G and the current global temperature is t , then

$$\text{Forecast} = \text{LW}(G) + \frac{1}{2}t.$$

But if it is Right's turn to play from G and the current global temperature is t , then

$$\text{Forecast} = \text{RW}(G) - \frac{1}{2}t.$$

We note that if the players both follow Sentestrat, then the forecast does not change during any interval of ko-free play at which the temperature remains constant. But, when Left moves from G , of temperature t_{old} , to G^L , of temperature t_{new} , the forecast changes from

$$\text{LW}(G) + \frac{1}{2}t_{\text{old}} \quad \text{to} \quad \text{RW}(G^L) - \frac{1}{2}t_{\text{new}}.$$

If Left followed Sentestrat, then $\text{RW}(G^L) = \text{LW}(G) + t_{\text{old}}$, and the effect of Left's move was not only to lower the temperature from t_{old} to t_{new} , but also to increase the forecast by

$$\Delta\text{Forecast} = \frac{1}{2}|\Delta t|.$$

In general, whenever the temperature drops by Δt , the forecast may be revised by up to $\frac{1}{2}|\Delta t|$ in favor of the player who made the last move at the old temperature.

Since either player may choose not to follow Sentestrat, other moves may occur that cause the economist to revise his forecast. However, it is not hard to see that all such forecast revisions are adverse to the player who moves. Hence, moves that cause forecast revisions adverse to the mover are called *sacrificial*. From this viewpoint, all moves can be partitioned into three sets:

1. Temperature-lowering moves improve the forecast by up to $\frac{1}{2}|\Delta t|$.
2. Sacrificial moves degrade the forecast.
3. All other moves leave the forecast unchanged.

Sacrificial moves may be either voluntary or compulsory, depending on whether or not any alternative option was available. An examination of Sentestrat reveals the only condition under which a sacrifice can be compulsory:

4. A sacrifice is voluntary *unless* the only move at the current temperature is banned by the ko rule.

So far, we have assumed that our economic kibitzer is able to compute a single thermograph for the entire game. In practice, this assumption may entail too much computation to be practical. In particular, if the game is a sum of many components, then it may be feasible to compute the thermograph of each component even though it may not be feasible to compute the thermograph of the sum. Under such circumstances, our economic kibitzer may be forced to define temperature and make forecasts based only on the component thermographs. And the resulting temperature may be higher. In particular, if two components sum to zero, then their thermographs will have the same temperature, t , and this will be the temperature used by the kibitzer who sees only the thermographs of the individual component games. However, the better-informed kibitzer who also knows the thermograph of the combined game will see it to be a vertical mast of minimum temperature.

Nevertheless, we assert that our economic guru can still classify each move that changes the forecast as temperature-lowering or sacrificial, even though forecasters using different databases may take somewhat differing views of the same game.

And it is easy to see that as the game approaches its conclusion, the forecast converges to the final score.

Evidently, the forecast is a plausible prediction of the final score. Thermostat is a greedy algorithm, which optimizes this forecast on a one-move time horizon.

In order to be a viable investment, our sacrificial move must prepare for some later payoff, and this payoff must eventually show up either as a temperature at which we get the last big move before a sizable drop in temperature, or as a ko that we can win while the opponent has only smaller (i.e., cooler) moves available.

If there are no foreseeable kos, then all changes in forecasts result from a sequence of plays in which the player who benefits from the forecast revision “gets the last big move”. Since each such event lowers the temperature, it is clear that the forecast’s absolute margin of error can be no worse than $t/2$. In very simple situations, the forecasting error is often precisely that. However, in a complicated game that is the sum of many components, it is reasonable to expect that the temperature will decline adiabatically in many small decrements, and that each side will be able to get about half of the total absolute value of these decrements. Under such conditions, the net of forecasting improvements and degradations will be about zero. We now concoct some models in which this result is precise and provable.

4.3. Enriched Environments. Let 2δ be the greatest common divisor of all temperatures of all positions of all summands of G . For example, if these temperatures are $\frac{1}{4}$, $\frac{1}{3}$, and $\frac{5}{8}$, then $2\delta = \frac{1}{24}$. Suppose the maximum initial

temperature of all summands of G is T . Then we define the *minimally enriched environment* MEE for G as the sum of the elementary switches of temperatures $\delta, 2.\delta, 3.\delta, 4.\delta, \dots, n.\delta$, where $n = T/\delta$. Let h be an upper bound on the total number of moves, in all positions of all summands of G , which might serve as a kothreat for any ko in G . (Any move that raises the temperature might serve as a kothreat). We then define SEE, the *standard enriched environment* for G , as $(2h + 1)$ copies of MEE. SEE is the sum of $(2h + 1).n$ distinct switches. An *over-enriched environment* is defined in the same way as the standard enriched environment, except that δ is taken to be some proper divisor of the greatest common divisor of all temperatures of all positions of all summands of G ; the maximum temperature of OEE may be larger than the temperature of G , and the number of switches at each multiple of δ may be any odd number greater than $(2h - 1)$.

Theorem on Enriched Environments. Let G be a position that contains no hyperactive kos. Then:

THEOREM. *The economic forecasts for $G + \text{MEE}$ and $G + \text{SEE}$ are the same as the economic forecast for G . If $G + \text{SEE}$ is played by two gurus using Conventional Rules, the economic forecast of the final score will be precisely correct.*

If our economist provides us with all of the relevant means and temperatures, we can play $G + \text{SEE}$ against an opposing guru. Even in the conventional game, we can elect to play Sente strat, subject to the proviso that whenever any move permitted by Sente strat involves a ko, we will prefer the ko move, and if there are several permitted ko moves, we will play the same one (if any) that we previously played most recently.

It is not hard to see that this policy ensures that the overall result is at least as favorable as the forecast. If a placid ko has temperature t , then the standard enriched environment also contains a bountiful supply of switches at the same temperature. Our policy ensures that the ko will be resolved before all of these switches are played. So, whenever we are faced with a koban, we have another (switch) move at the same temperature that serves us just as well.

A Sequence of Universal Environments. Notice that it is possible to construct a sequence of over-enriched environments that are universal. We'll construct a specific sequence, U_1, U_2, U_3, \dots , with the property that if G is *any* game that contains no hyperactive kos, then for all sufficiently large n ,

- the Economist's forecast for $G + U_n$ is the same as the forecast for G , and
- the forecast of $G + U_n$ is precisely accurate.

To this end, it is sufficient to let

$$\delta_n = \frac{1}{\text{lcm}\{1, 2, 3, \dots, n\}}.$$

Let U_n consist of $(2n + 1)n/\delta_n$ switches, $(2n + 1)$ copies each of temperatures δ , 2δ , \dots , n .

It is easy to see that this construction has the specified property; it is a sequence of “Universal Enriched Environments”. Any reasonable game G (one without hyperactive kos), when played in such an environment, behaves precisely according to the economist’s forecast.

Universal environments, and other enriched environments, are artificial theoretical constructions intended to show why, when there are sufficiently many independent regions in the game, the economic forecast is a very plausible prediction of the outcome.

In fact, there are serious difficulties in realizing our theoretical enriched environments within the practical constraints of a game like Go, even if we allow it to be played as the sum of arbitrarily many boards, some of which may have arbitrarily large sizes. The first problem is that all final scores in Go are integers, and this Diophantine constraint prevents the direct construction of even simple switches, such as $\{\frac{1}{2} | -\frac{1}{2}\}$. This difficulty can be at least partially alleviated by working with “chilled Go”, as described in [Berlekamp and Wolfe 1994]. But even in chilled Go, there is no switch of value such as $\{\frac{1}{3} | -\frac{1}{3}\}$; the denominators of all temperatures of switches must be powers of 2.

So, even though the enriched environments we have constructed are all sums of extremely simple (two-stop) abstract games, some of these games might not be realizable within the natural constraints of the class of games (e.g., Go) to which we want to apply our economic forecasts.

In more realistic environments, our economic forecasts are typically not precise. But they are still often quite close. Yonghoan [Kim 1995; Berlekamp and Kim 1996] has analyzed placid kos in several kinds of realistic environments, in which there are no other regions having exactly the same temperature as the ko. In such situations, kothreats have value, but the typical value of a kothreat is only a small multiple of δ , where (as in our constructions), δ is the typical difference between temperatures of other regions.

So, the total value of any fixed number of kothreats still approaches zero as the background becomes sufficiently rich. The rich environment must have many independent regions, whose temperatures are relatively dense in the vicinity of the temperature of the placid ko. Other details about the games in the background environment don’t seem to matter much; they need not be switches. Even though the details of the background can affect the accuracy of the forecast by up to $t/2$, nearly all of these forecasting inaccuracies are due to inherent difficulties in predicting who can get the last move at each tax rate. Only a small part of the forecasting inaccuracies are due to placid kos, even though thermographs ignore kothreats entirely.

Hyperactive Kos. Even in a very rich environment, kothreats can have a big effect on the final score if the game contains a position that is a hyperactive ko.

Buyer	Item	Amt Bid	Book Value	Max \$ Loss		Max Chip Loss	
				DW	YK	DW	YK
YK	Black in -Go	$2\frac{1}{4}$	$2\frac{1}{32}$		$\frac{7}{32}$		
DW	Black in -Go	$1\frac{3}{4}$	$2\frac{1}{32}$	$\frac{9}{32}$			
YK	Domin. Left	$1\frac{1}{2}$	$\frac{3}{4}$		$\frac{3}{4}$		
DW	Domin. Left	$\frac{1}{2}$	$\frac{3}{4}$	$\frac{1}{4}$			
DW	First move	$1\frac{3}{4}$	$1\frac{15}{32}$	$\frac{9}{32}$		$\frac{9}{64}$	
YK	First move	$1\frac{1}{4}$	$1\frac{15}{32}$		$\frac{7}{32}$		$\frac{7}{64}$
Total mistakes				$\frac{43}{64}$	$\frac{69}{64}$	$\frac{9}{64}$	$\frac{7}{64}$
Net gains due to bidding				$+\frac{13}{64}$	$-\frac{13}{64}$	$-\frac{1}{64}$	$+\frac{1}{64}$

Table 2. Accountant's error ledger for the demonstration game of Figures 2-6.

From the economist's perspective, the μ -value of such a ko depends on who is its master, and this, in turn, depends on a global tally of kothreats. And, en route to elegance and simplicity, combinatorial game theory suppresses much kothreat information from the canonical form. Then even more kothreat information is suppressed when the information in the canonical form is condensed into a thermograph. So, if there is a hyperactive ko, a precise analysis requires more data than the mathematics has retained. Of course, $\mu(\hat{G})$ and $\mu(\check{G})$ still provide firm bounds; but for some positions these bounds may be quite far apart.

4.4. Dollar Accounting. Economic Rules facilitate some very detailed accounting, which identifies the moves at which mistakes were made and the cost of each mistake. Table 2 shows an example.

In these auctions, every Book bid was intermediate between the bids submitted by the two players, and all maximum bidding losses were realized.

Altogether, in these auctions, the average of the two players made 1 point worth of bidding errors, $\frac{7}{8}$ in dollars, and $\frac{1}{8}$ in chips. DW was better than this average; YK was worse. The absolute difference of each from the average was $\frac{13}{64}$ dollars minus $\frac{1}{64}$ chips.

After all play was concluded, the final result (after converting chips back to cash) was that DW ended with \$5.1875; YK, with \$4.8125. Since both players played all moves according to perfect Book strategies, YK's net loss of $\frac{3}{16}$ resulted entirely from the fact that his bidding errors cost more than DW's.

The economist's accounting can be readily extended to itemize any errors that occur in play.

DW	YK		Total chips still on the table	Temper- ature
-2G - 4Y	-2G - 4Y	At start	4G + 8Y	$1\frac{1}{2}$ (bid)
1G + 2Y		(after move 1)	3G + 6Y	$1\frac{1}{8}$
1Y		(after move 5)	3G + 5Y	$1\frac{1}{16}$
	1Y	(after move 6)	3G + 4Y	1
2Y		(after move 9)	3G + 2Y	$\frac{7}{8}$
	2Y	(after move 10)	3G	$\frac{3}{4}$
1G		(after move 11)	2G	$\frac{1}{2}$
	2G	(after move 12)	0	0
+ Y	- Y			
= + $\frac{1}{16}$	- $\frac{1}{16}$	Net total chips gained from playing (both players used perfect Book play)		

Table 3. Pocketed chip ledger for the demonstration game of Figures 2–6.

4.5. Chip Accounting. The accuracy of economic forecasts is limited by a crucial approximation that states that, in the course of a sufficiently long and intricate well-played game, each player will eventually pocket about half of the chips that are placed on the table at the conclusion of the initial auctions.

It is instructive to see how the forecasts evolved during the (unenriched) demonstration game. At the start of the game, each player put chips of value $\frac{3}{4}$ onto the table ($Y = \frac{1}{4}G = \frac{1}{16}$). This total sum of $1\frac{1}{2}$ units ($4G + 8Y$) was then pocketed as shown in Table 3.

4.6. Forecasting Records. Economists studying history like the Big Picture, which, for the demonstration game, is shown in Figure 22. As usual, μ is plotted along the reversed horizontal axis; t , on the vertical. Each point plotted in this figure shows the global $[\mu, t]$ values of the full game after the numbered move.

Figure 22 records the values of the total μ and the maximum t after moves 0–12 of the demonstration game. The sequence terminates when the temperature reaches 0. The values of μ oscillate back and forth around the forecast, and each decrease in the tax rate is accompanied by a revision in the forecast, as illustrated in Table 4.

The magnitude of the revision is one half the decrease in the tax rate, and the sign of the revision favors the player who made the last move before the tax rate decrease.




Figure 22. The Big Picture of the demonstration game.

Tax Rate	μ	Forecast	Move	Change
$1 \frac{15}{32}$	$2 \frac{25}{32}$	$2 \frac{3}{64}$	1	$-\frac{11}{64}$
$1 \frac{1}{8}$	$1 \frac{5}{16}$	$1 \frac{7}{8}$	2-5	$-\frac{1}{32}$
$1 \frac{1}{16}$	$1 \frac{5}{16}$	$1 \frac{27}{32}$	6	$+\frac{1}{32}$
1	$2 \frac{3}{8}$	$1 \frac{7}{8}$	7-9	$-\frac{1}{16}$
$\frac{7}{8}$	$1 \frac{3}{8}$	$1 \frac{13}{16}$	10	$+\frac{1}{16}$
$\frac{3}{4}$	$2 \frac{1}{4}$	$1 \frac{7}{8}$	11	$-\frac{1}{8}$
$\frac{1}{2}$	$1 \frac{1}{2}$	$1 \frac{3}{4}$	12	$+\frac{1}{4}$
0	2	2	13 ...	

Table 4. Forecasting record.

4.7. The Komi in Go. Professional Go games begin with an empty board. By symmetry, $\mu = 0$. The player who makes the first move is required to give his opponent some number of points in return. This number of points is called the *komi*. To prevent ties, the komi is traditionally chosen as a half-integer. The komi is specified by the tournament directors. Values of the komi in modern tournaments have ranged from 5.5 to 8.5 points.

Our economist's viewpoint suggests that the komi could be determined by a competitive auction between the two players rather than by executive decree. Our economist conjectures that the optimal value of the komi, like his forecast, should be half the temperature of the empty board.

If the temperature of the empty board is t , then an expert player will find the first move on the empty board to be equally attractive to playing on an additive switch of value $+t|-t$. In Go jargon, such a switch is called a *2t-point gote move*. So, from the viewpoint of our economist, the n -point gote move that is just as tempting as the first move on an empty board is attained when n is *four* times the value of the fair komi. That's because n should be twice the temperature, and the fair komi should be half of the temperature. In view of the accepted range of the komi, n is presumably twenty-something.

5. Incentives

In loopfree combinatorial games without taxes, the best moves are those for which the mover has a maximum incentive. If the position is G , Left's incentive is $\Delta^L(G) = G^L - G$; Right's incentive is $\Delta^R(G) = G - G^R$. If $G = A + B + C + \dots + F$, then $\Delta^L(G)$ and $\Delta^R(G)$ are the sets

$$\begin{aligned}\Delta^L(G) &= \{\Delta^L(A), \Delta^L(B), \Delta^L(C), \dots, \Delta^L(F)\} \\ \Delta^R(G) &= \{\Delta^R(A), \Delta^R(B), \Delta^R(C), \dots, \Delta^R(F)\}\end{aligned}$$

It is helpful to consider the union of both sets:

$$\Delta = \{\Delta^L, \Delta^R\}.$$

When temperatures are negative, both incentives equal the temperature, and values are numbers that are equal to the means. However, when temperatures are nonnegative, both values and incentives are necessarily games that are more complicated than numbers. To manipulate these games, familiarity with classical combinatorial game theory is required [Berlekamp et al. 1982; Conway 1976].

The partial ordering of the incentives plays an important role in the mathematical analysis of various lines of play [Berlekamp and Kim 1996]. If $\Delta^L(A) > \Delta^L(B)$, Left will prefer to play on A rather than on B. If $\Delta^R(A) > \Delta^R(B)$, then Left's move on B is reversed through Right's response on A. Reversibility is a technical criterion well known in combinatorial game theory.

Every incentive is a game that has 0 as a right follower. Conversely, any game in which 0 is a right follower is one of its own right incentives. So no incentive can be 0 or positive.

Many pairs of incentives are incomparable. On the other hand, temperatures of incentives, being numbers, can be totally ordered (although ties will occur in cases of equal temperatures). Thus, a wise prelude to an investigation of the partial ordering of the incentives is to compute each incentive's temperature, and then sort them accordingly. If $\text{temp}(\Delta^L(A)) > \text{temp}(\Delta^L(B))$, this does

not imply that $\Delta^L(A) > \Delta^L(B)$, but it certainly precludes the possibility that $\Delta^L(B) \geq \Delta^L(A)$.

In general, $\text{temp}(A - B) \leq \max\{\text{temp}(A), \text{temp}(B)\}$. However, in the special case of incentives, we can assert that for any G , there is at least one left incentive such that

$$\text{temp } \Delta^L(G) = \max \{ \text{temp}(G), \text{temp}(G^L) \}.$$

This is true because if the temperature of the incentive were less, then we could pick a value of t between the temperature of the incentive and the temperature of G and G^L . Cooling the incentive-defining equation by t would yield

$$\Delta^L(G_t) = G_t^L - G_t,$$

where neither term on the right side is a number, but their difference is. But this contradicts another theorem of combinatorial game theory, which states that if $H = G_t$ is not a number, then it has at least one left incentive that exceeds all negative numbers.

The same arguments can be repeated for right incentives.

In earlier sections of this paper, we showed that for economic rules, when players bid to set a new tax rate, either player can safely select the temperature of the full board position as his bid, except in the case when he has no legal move. This leads directly to the important conclusion that initial tax rate bids may be submitted independently of the bids to play Left's side of the game. When the tax rate is fair and both players are competent, it makes no economic difference who moves first.

However, it is in many ways more natural to view the temperature as a function of the move rather than as a function of the initial position. In that sense, temperature is a natural function of the incentive. Sente strat still works.

5.1. "Cooling" via Honorific Tie-Breaking and Prescribed Initial Temperatures. The winner of an economic game is the player who realizes a net profit. His opponent necessarily incurs a net loss. If two gurus bid and play perfectly, then the net score should be zero and the economic outcome is a tie. We can break such ties by awarding an honorific victory to the player whose opponent made the first pass at the original temperature at which the game started.

We can also decree that the economic game begin with a prescribed finite tax rate, t , and with a prescribed player assigned to play first.

We temporarily assume that the sum is loopfree. If t is sufficiently large, this game is identical to the rules of our ideal economist, augmented only by the honorific tie-breaking convention. However, if t is 0, then this game is identical to a classical combinatorial game with the last mover winning. This is because at subzero temperatures, the economic rules and the conventional rules yield identical outcomes.

At intermediate values of the starting temperature, this game turns out to be identical to G_t , called “G cooled by t ”. Cooling is a well-known mapping of combinatorial games onto combinatorial games.

5.2. Implications for Computer Go and Similar Games We focus on those endgames that are sums of numerous regional positions. Such problems best display the unique power of mathematics to cope with complexities that lie beyond the reach of alternative methods based only on heuristics and/or the judgment of human experts.

Combinatorial game theory provides the machinery to obtain precise analyses of sums of loopfree games played according to conventional rules. The recommended approach is to compute the incentives of all summands, then order the incentives as much as possible, and then search among all lines of play that are locally canonical. As exemplified by the \$1000 Ko problem [Berlekamp and Kim 1996], this approach often yields a rigorous analysis relatively quickly, even for problems that are well beyond the scope of less sophisticated methods. However, it is known that, in general, even sums of simple games can be NP hard. The first such results were due to Lockwood Morris. Later, Laura Yedwab [1985] and David Moews [1993] showed that a precise determination of the outcome of even a sum of elementary three-stop games is NP-hard. So, for some problems in today’s technology, the preferred methodology of combinatorial game theory (like any other method that finds only rigorously correct answers) must necessarily bog down in a combinatorial explosion of calculations that will not terminate until long after everyone has run out of patience.

On the other hand, thermography offers fast, polynomial-time algorithms that solve the same problems according to slightly different (economic) rules. The discrepancy between the ideal results of sums of games played according to economic rules and conventional rules is relatively small and easily bounded. Economic forecasting offers further small adjustments that usually yield further decreases in this discrepancy, driving it towards zero as the background environment included in the sum becomes sufficiently rich. However, in the most difficult problems, the background may still be “impoverished” despite the presence of many summands.

We have given an economic interpretation of the combinatorial game theory technique of *cooling*. This interpretation suggests a strategy by which computer Go buffs might achieve appropriate tradeoffs that would combine most of the precision of incentives with most of the speed of thermography. This approach begins by computing thermographs of positions and their incentives, top-down. This provides a solution of the frozen image of the original game. Then one refines these conclusions to obtain a solution to the original game cooled by a temperature that is not quite large enough to freeze it. This can then be further refined by further iterations, each of which yields a solution of a less-cooled version of the original.

5.3. Thermal Disassociation. Thermal disassociation [Berlekamp et al. 1982], first proposed by Simon Norton in the early 1970s, represents an arbitrary loopfree game as the sum of its mean and appropriate infinitesimals, each heated by a successively smaller temperature. Excellent approximations to the game, at various temperatures, can be obtained by ignoring the cooler terms in this expansion.

A ko can also be represented as the sum of its mean and a heated “infinitesimal ko”. The study of infinitesimal kos is still in its infancy.

As a loopfree instructive example, we consider the incentives of the general 3-stop game, $G = x|y||0$. G’s stops are 0 and the numbers x and y , where $x \geq y \geq 0$. G’s right incentive is $\Delta^R(G) = G - G^R = G$. Its left incentive is

$$\begin{aligned} \Delta^L(G) &= G^L - G \\ &= \{x|y\} + \{0||-y|-x\} \\ &= \{x||x-y|0, x|y||0\} \\ &= \{x-y, x|y||0\} \\ &= H. \end{aligned}$$

As shown in Figure 23, the thermographs of H can have any of three different forms, depending on how x compares with $2y$ and $3y$.

Corresponding to Figure 23, we have these expressions for the thermal dissociations of G and H, where $\neg_v = \{v|0||0\}$ and $\dagger_v = \{0||0|-v\}$:

Cases 1 and 2: $G = \frac{x+y}{4} + \int^{\frac{1}{4}(x+y)} * + \int^{\frac{1}{2}(x-y)} \varepsilon_1,$

where $\varepsilon_1 = \{*\|\frac{1}{2}(x-3y)\} + \{\frac{1}{2}(3y-x)|0\}$ is fuzzy with 0

Case 3: $G = y + \int^y \neg_{x-3y}$

Case 1: $H = G$

Case 2: $H = G + \int^{x-2y} \varepsilon_2,$

where $\varepsilon_2 = \{0, 3y-x|0||x-3y\} - \{3y-x|0||x-3y\} > 0$

Case 3: $H = \frac{x-y}{2} + \int^{\frac{1}{2}(x-y)} * + \int^y \dagger_{x-3y}$

(Note: We have $H > G$, since $v|0 > \neg_v \neg_v$.)

It is interesting to note that the Yedwab–Moews construction of NP-completeness requires no games of Case 2, and only their coolest game is of Case 3. All but one of their summands are of Case 1, in which left and right incentives are equal.

$$G = \{x|y||0\} = \Delta^R(G)$$

$$H = \{x-y, x|y||0\} = \Delta^L(G)$$




Figure 23. Thermographs of incentives of three-steps.

References

[Berlekamp and Kim 1996] Elwyn R. Berlekamp and Yonghoan Kim, “Where is the thousand dollar ko?”, pp. 203–226 in this volume.

[Berlekamp and Wolfe 1994] Elwyn R. Berlekamp and David Wolfe, *Mathematical Go: Chilling Gets the Last Point*, A. K. Peters, Wellesley, MA, 1994. Also published in paperback, with accompanying software, as *Mathematical Go: Nightmares for the Professional Go Player*, by Ishi Press International, San Jose, CA.

[Berlekamp et al. 1982] Elwyn R. Berlekamp, John H. Conway, and Richard K. Guy, *Winning Ways for Your Mathematical Plays*, Academic Press, London, 1982.

[Conway 1976] John H. Conway, *On Numbers and Games*, Academic Press, London, 1976.

- [Hanner 1959] Olof Hanner, "Mean play of sums of positional games", *Pacific J. Math.* **9** (1959), 81–99.
- [Kim 1995] Yonghoan Kim, "New values in Domineering and loopy games in Go", Ph.D. thesis, Mathematics Department, University of California at Berkeley, May 1995.
- [Milnor 1953] John Milnor, "Sums of positional games", pp. 291–301 in *Contributions to the Theory of Games*, vol. 2 (edited by H. W. Kuhn and A. W. Tucker), Ann. of Math. Stud. **28**, Princeton University Press, Princeton, 1953.
- [Moews 1993] David Moews, "On some combinatorial games connected with Go", Ph.D. thesis, Mathematics Department, University of California at Berkeley, December 1993.
- [Yedwab 1985] Laura J. Yedwab, "On playing well in a sum of games", M.Sc. Thesis, MIT, 1985. Issued as MIT/LCS/TR-348, MIT Laboratory for Computer Science, Cambridge, MA.

ELWYN BERLEKAMP
DEPARTMENT OF MATHEMATICS
UNIVERSITY OF CALIFORNIA AT BERKELEY
BERKELEY, CA 94720
berlek@math.berkeley.edu