

This synopsis is for the use of ACBL Members to promote greater understanding about Computer Based Cheating Detection by the ACBL. What follows are excerpts from existing public reports. The ACBL, the EDGAR Foundation, and the Institute for Bridge Arbitration are responsible for their respective activities. See [ACBL Computer Based Cheating Detection](#) for details and sources.

Twenty months of testing preceded first pilot EDGAR runs (Jan 2024). More than 35000 online ACBL players have been assessed. Edgar requires large amounts of evidence to identify a cheating player. Skill level, propensity to make wild plays, unexpected success all have no bearing on how EDGAR separates cheaters from normal players. Thorough testing with actual data from known cheaters and general players has vetted the tool's accuracy and reliability.

"The system is comprised of several automated "detectors", each focusing on one aspect of the game. All detectors, whether the focus is bidding, opening leads, or defense, are constructed similarly: a player or partnership has an opportunity to do something that looks incriminating (a "hit"), absolving (a "miss"), or neutral. Depending on the detector, hit/miss might be nothing more than a good or bad decision/guess. One hit or miss only moves the needle a little bit, so the results are aggregated across many plays and many detectors before reaching a conclusion with confidence. How many deals depends on how blatant the cheating."

"EDGAR looks for the run-of-the-mill decisions which an honest pair is going to get wrong some of the time, not just the "smoking gun" ace underleads which hit partner's king-doubleton. If a pair gets an unusually low number of them wrong, that is an indication that the pair may be cheating."

"EDGAR does not consider any one action suspicious. There is no play, no matter how extreme or profitable, that will trigger a response from EDGAR. Your opponents could be apoplectic with rage over that unbelievable play you just made, yet EDGAR will not notice it until you've done it over and over with a highly abnormal rate of success."

"EDGAR is a system for weighing evidence, both incriminating and absolving, in many facets of the game. It notices when offenders make lots of consistently successful weird leads, but the data it collects also exonerates honest players whose playing style is to take lots of unusual actions that sometimes strike gold. It doesn't matter whether they do this from inexperience, or simply because they have an idiosyncratic playing style. An occasional run of luck by an innocent player might upset the opponents, but it won't upset EDGAR. Their many weird leads will simply speed their reversion to the mean."

DETECTORS

Leads

"Leads away from an unsupported King

Opening Lead Detector (KOM) = retains the objectivity and tirelessness of double-dummy, but conforms much more to expert analysis than double-dummy does. ...some hits and misses are more meaningful than others. In this, it follows a more traditional investigation approach, which also typically distinguishes the strength of evidence in addition to classifying it as incriminating or absolving. For KOM, the more "swinging" a lead is, and the more it out- or under-performs plausible alternatives, the more points won or lost.

Opening Leader's double-dummy (DD) error rate is another useful tool for detecting cheating.

(1) An Investigator evaluates the opening lead on a single-dummy (SD) basis, only considering Opening Leader's cards and the auction.

(2) A judge compares the competing leads (the actual lead vs the Investigator's lead) looking at what a colluder can see, i.e., the cards held by both defenders in that suit. We call this "partner-dummy" (PD), as it is neither single-dummy nor DD.

(3) Results are tallied and interpreted.

KOM gives more weight to some wins and losses than others. The more "swingy" a lead is, and the more it out- or under-performs plausible alternatives, the more points won or lost.

KOM has separate detectors for OLs vs suits and OLs vs NT.

Other OL detectors look at specific situations:

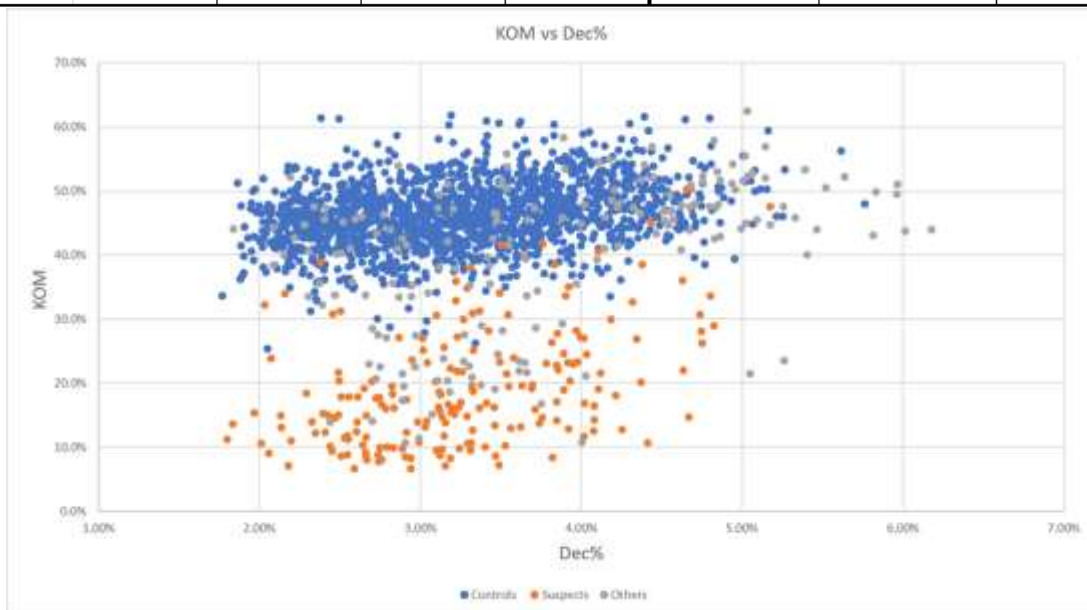
(1) **Weird Leads** is a separate detector for the kinds of leads that experienced players almost never make, such as leading low from AQxx or KQJx against a suit contract.

(2) **Trump Leads** are also excluded from KOM. Evaluating the attractiveness of a trump lead is more complex than other suits (or at least we haven't solved it well enough yet). We still note their frequency and evaluate them in other ways.

(3) **Confusing Auctions**. Edgar records explanations of calls, but players don't include them consistently and don't intend to undertake the sort of difficult natural language processing that would be required to understand them. However, we have put a fair amount of work into understanding most auctions reasonably well (not as well as a human, of course). Very occasionally, though, EDGAR throws up its virtual hands and says it can't figure out the auction. Such deals are excluded from KOM (we keep count in case it's unusually common for a pair).

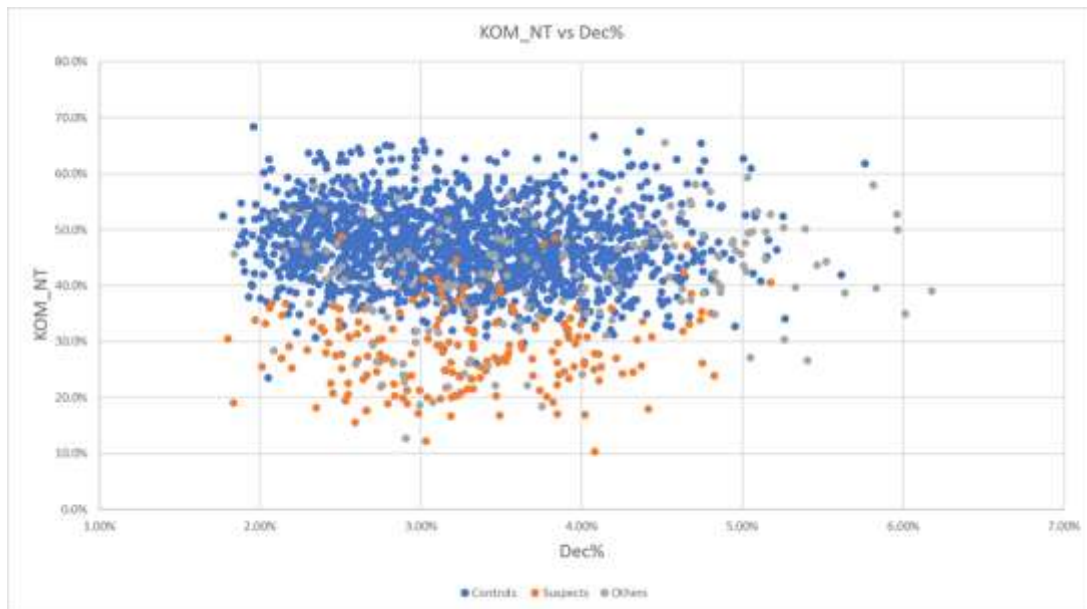
(4) **Leads against slams."**

Group	Total Players	Dec%	LDL%	BL%	KOM miss rate	KOM_NT miss rate	DD OL miss rate
Suspects	193	3.3%	8.5%	2.7%	19.5%	27.0%	15.7%
Controls	1596	3.3%	11.7%	6.0%	46.3%	47.4%	20.6%
Others	176	3.8%	11.5%	5.1%	40.1%	41.9%	20.0%



Control players are in blue – the suspects are in Orange.

The better the separation of the cloud the better the precision of the signal.



Weird Leads/Plays and Exonerating the Inexperienced Player

“The first table shows the number of weird leads per 1000 opening leads and the miss rate among groups sorted by playing ability (DEC%, which is double dummy error rate when declaring, is used here as a simple proxy for ability). Only pairs not flagged by EDGAR are included. We can see that the frequency of weird leads rises dramatically as playing ability drops, ranging from rare among experts to commonplace among novices. Curiously, the success of weird leads is nearly independent of ability level. Despite whatever inferences are available to them, experts appear to be no better than novices at hitting their partner’s touching honor when deciding to underlead these holdings.

Table 1

	Playing Ability					
	BETTER					WORSE
DEC%	< 3%	3-3.5%	3.5-4%	4-4.5%	4.5-5%	> 5%
WL/1000	0.6	1.5	3.7	8.4	13.6	25.4
Miss Rate	54%	52%	52%	53%	52%	52%

The second table shows the same type of results for pairs flagged by EDGAR as cheating, which is 3% of all hands in this dataset:

Table 2

	Playing Ability					
	BETTER					WORSE
DEC%	< 3%	3-3.5%	3.5-4%	4-4.5%	4.5-5%	> 5%
WL/1000	4.0	5.4	11.3	18.9	29.3	35.5
Miss Rate	9%	10%	15%	21%	23%	33%

Two differences from the previous table stand out: cheaters make weird leads much more often than non-cheaters, and they miss much less frequently doing it. The biggest improvement is found among experts, but


even the most inexperienced players achieve a sizable reduction in miss rate. (See Appendix I for a more detailed breakdown of the numbers.)

The Weird Lead behavior of cheating players is very different from that of their non-cheating peers.

	WL/1000	Miss Rate
Cheating Partnership	18.4	14%
With Other Partners	4.4	51%

It is easier to understand these things by looking at hits and misses instead of miss rates. What we see in figure 1 (below) is that the cheating cohorts make unsuccessful weird leads almost as often as the non-cheating cohorts, but, at every skill level, they make far more successful weird leads.

Figure 1



At the expert end of the spectrum, there are so few unsuccessful weird leads that the overall hit-rate for the cheating experts is very high. At the other end, the hit-rate is weighed down by the many unsuccessful leads still made, though the number of “excess” hits is actually much higher.

We think there are a number of reasons for this:

- Inexperienced players may lead “4th best” even from 3 honors just because they think it’s the right lead.
- Not every cheater cheats perfectly, or all the time

- Very rarely a player may be “smarter than EDGAR” and see that not catching a touching card will still work out fine.”

Bidding

Bad Landings – ...”playing in a 4-2 or 3-3 or worse fit; Stopping in 5M when slam is cold, Playing in a minor suit when a superior major fit is available, bidding over opponent’s preempts, opening or overcalling one spade then finding an 8-card heart fit; Finding a better minor after opening a Precision 1♦ or a 2+ 1♣ in 2/1 GFR. Many different unrelated auctions funnel to a similar end point. No bridge player gets every hand right; Bad Landings was developed with the assumption that colluders would significantly outperform clean players in “guessing” well, and avoiding accidents.

- First, as expected, the Cheaters make 40% more DD errors when declaring than the Top Clean players. Without going into detail here, this is a material skill difference.
- Second, the Top Cheaters have almost half the Bad Landings as the Top Clean players.
- In other words, bridge skill can only do so much to eliminate Bad Landings, and the effects of cheating stand out in comparison.
- (It is also interesting that the BL% (Bad Landing percentage) for Top Clean Players is not so different from All Clean Players.)”

Defense

Defensive cardplay, called LaterDefLeads (LDL for short)

“LaterDefLeads (LDL) counts the double-dummy errors by each defender, but only for tricks 2-12 and only when a defender leads to the trick. It excludes the opening lead, and any cards defenders play when following suit or discarding.

Score = % of trick 2-12 leads by defenders which are double-dummy errors –

Looking at a dataset of 600 random pairs who each played over 1000 deals, defender’s card plays, excluding the opening lead and trick 13 were:

86% following suit + discards

14% leading to a trick

Even though only 1/7 of defender’s cards are LaterDefLeads, they accounted for 60% of their DD errors. Following suit & discards were roughly 20% each.”

... “colluding would benefit defenders the most when on lead, as opposed to following suit or discarding (which we believe are more prone to concentration errors, and so less useful for distinguishing collusion). ... We expect better players to make fewer errors, so skill level should be accounted for. Here are the same players, but the LDL% is plotted vs Dec% (DD Declarer error rate). Plotting LDL% by Dec% greatly reduces the overlap between the (Controls and Suspects). ... Because it varies with a skill estimate, we took a conservative approach to setting the expert number. The intent is to reduce the chance of false positives stemming from world-class play or eccentric style.

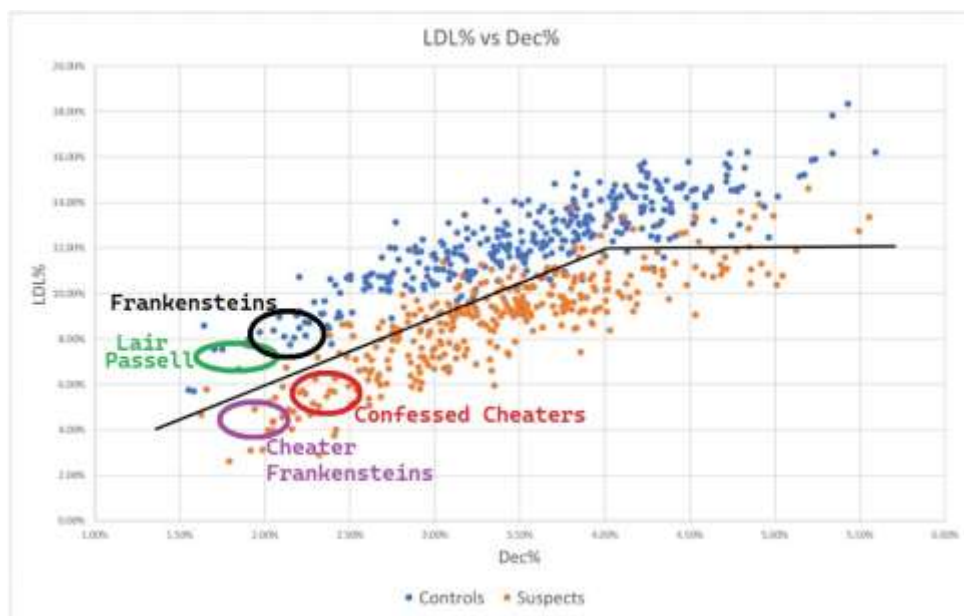
Can EDGAR differentiate between winners and cheaters?”

Yes – EDGAR created two “pairs” from aggregating games from all players who won a large number of ACBL events, and all players who finished last, or lost, these events. These artificial “pairs” compare against known cheaters and their performance. Known cheaters come from the ACBL discipline list. Since these “pairs” are artificial so the authors named them Frankensteins – there are four:

	Frankenstein	Loser Frankenstein		Winning Cheater Frankenstein	Losing Cheater Frankenstein
Dec%	1.9%	2.3%	Dec%	1.8%	2.1%
LDL	7.5%	8.8%	LDL	4.0%	4.9%
KOM	43.7%	45.2%	KOM	9.0%	14.2%
BL	4.8%	5.5%	BL	2.7%	2.4%

	Non-cheating Winners	Cheating Losers
Dec%	1.9%	2.1%
LDL	7.5%	4.9%
KOM	43.7%	14.2%
BL	4.8%	2.4%

Here Dec% is the Double Dummy Declarer Score – a measure of skill (lower is better) not cheating per se. LDL is Later Defensive Leads – double dummy defending scores; KOM is the opening leads score; and BL is Bad Landings.



EDGAR REPORTS – What ACBL Sends Accused Members

When a member is contacted by the ACBL for cheating online, they receive all EDGAR hands and data. This can be an overwhelming amount of information. You have a short period of time to respond (30 days). You may request an extension.

“The form Notice of Discipline informs you that enclosed with the notice are several documents including,

- 'Understanding EDGAR' which includes an overview and the reports that serve as the basis for ACBL's action against you"
- "[a] spreadsheet which includes each deal EDGAR analyzed."
- A "Concise EDGAR Report"

- document with "HOTSPOT REPORTS" as part of its title.
- The two-page "EDGAR Overview" refers to the latter document as "EDGAR Hotspot Reports" and offers explanations of the significance of each report."

..."(A) table that lists the performance of you and your partner on eight rows of EDGAR "detectors" that are listed in one column, sometimes giving each performance an EDGAR score in another column. Four of the detectors ("BadLandings," "PassedForce," "Bad_6m," and "RKC_5M") apply to the pair and the other four ("LaterDefLeads," "KOM_Suits," "KOM_NT2," and "Weird_OLs") apply to each individual partner, whose BBO names are listed.

The overview explains: " 'EDGAR points' start to accrue when play is materially better than normal expert play. . . [T]o amass enough EDGAR points to be charged under ACBL's CDR, a pair must exceed the 'expert' levels by a large enough margin over hundreds or thousands of deals, far beyond a plausible 'lucky streak'. A TOTAL EDGAR SCORE of 100+ is often the result of a pair using illicit information. Scores over 200 are exceptionally likely to indicate frequent influence from UI."

A large spreadsheet listing ALL HANDS surveyed by EDGAR is provided. The Information contained includes: "... (H)ere are the column headers in the complete data sets that are provided to people

Ref # (used for indexing)	Opener (Opener's hand)
Event name (date and time stamp for the event)	Responder
BBO (tiny url for the LIN file for the hand in question)	Direct-Seat
BN (per event index)	4-hand
N (player)	Con (What was the final contract)
S (player)	Tricks taken (How many tricks were taken)
W (player)	Max DD (What is the maximum number of DD tricks available)
E (player)	Dec (What seat declared)
Bidding (e.g. P-P-P-1S P-2S-P-4S P-P-P)	Dec name (Name of the declarer)
OB name (name of the opening bidder)	Dec Hand
Seat-Vul (number 1-4, and Val Info)	Dummy Hand
Opps bid	Leader (Name of the player on opening lead)
Doubled (did the final contract get doubled)	OL (What was the opening lead)
OB (What was the opening bid)	OL Suit
Direct (What happened in direct seat)	P Suit
Resp (What did responder do)	Leader Hand
4th (What did 4th seat do)	Third Hand
O-Rebid (What was Opener's rebid)	

There is no link between a Detector Score and a specific hand."

The hotspot report

"Having tried in vain to identify from the concise report and the spreadsheet a single board that is considered to be evidence of your misconduct, you turn to the multi-page hotspot report ("HSR"). What you find in this document is a series of ten reports collecting all actions (both bidding and leading) by you and your partner involving six hotspots. Two hotspots (PassedForce and Weird_OLs) are classified as partnership actions (even though the concise report lists Weird_OLs separately by partner) while the other four hotspots ("Suit_Overeasy," "Kxx_vsSuit," "Ax+_Low," and "SlamJumps") separately list your conduct and your partner's. You'll notice that the hotspot names are different than the detector names except for PassedForce and Weird_OLs.

Because the HSR lists all boards fitting six different patterns, depending on the frequency of your online play, the HSR may include 100, 200, 500, or 1,000 boards. The HSR is generous with the amount of information provided about each board listed under each hotspot, not only providing a hyperlink to the board's BBO movie, but also stating the board number, the date played, the contract, the opening lead, the auction, the hands held by you and partner, your BBO names, and whether the action has been labeled by EDGAR as a "hit," "miss," or "tie," sometimes even specifying the type of hit.

"EDGAR Hotspot Reports are intended to provide evidence similar to what might be presented during an IBA hearing." There is additional elaboration in the HSR introduction. Until it is updated, the introduction is likely to tell you: "This document contains a collection of 'hotspot reports' which are intended to highlight potential areas for further investigation. They are not part of the Concise EDGAR Report and play no role in EDGAR's Computer Based Cheating Detection, they are purely supplemental."

"Please note that the documents included with this notice are not necessarily the evidence that would be presented to an arbitration panel if you choose to challenge the charge"

EDGAR Reports are not evidence presented to the IBA in any arbitration. The ACBL decides what evidence to provide (shared during the preliminary hearing).

So, in deciding whether to opt for arbitration, despite having received pages of documentation including hundreds of boards, you cannot predict with any certainty what the outcome of arbitration might be. You have incomplete information and less than a month to decide."

Roles

It is important to understand that the ACBL has centralized handling of all online ethical violations. Clubs, Units, and Districts have no role nor any standing in their adjudication. The Managing Director, the Appeals & Charges Committee and the National Recorder execute the requirements in the ACBL Code of Disciplinary Regulations.

Process

If you receive a fast-track discipline notice from the ACBL, you can choose to accept the reduced penalty by completing the fair play pledge and ethics training online. The ACBL considers the effort to arbitrate a case substantial and wants to settle as many cases as possible prior to that step. IF you know you are innocent you can initiate arbitration to cause the ACBL to bring their case and give you the discovery information from which you will form your defense. Your case hearing occurs at an agreed time after the discovery hearing. The Institute for Bridge Arbitration has final say on innocence and on the severity of the punishment. Several claimants have had their penalties reduced substantially. Please see the current ACBL Code of Disciplinary Regulations for the current processes for all ethical violations.

Summary

I have compiled this synopsis to help answer many questions and doubts raised as the Computer Based Cheating Detection System touches players and clubs close to us. The EDGAR tool is robust and very accurate. EDGAR is set to avoid accusing innocent players, accepting that some cheating players will not get caught as a result. So far, no arbitration has exonerated any player, and wins have been limited to reduced penalties. The ACBL intends to scan everyone who has played online since March 2020. As of March 2025, 12 cases have been arbitrated from 784 accused players (from over 35,000 scanned). 18 cases are pending arbitration. The remaining players have either accepted their penalty or resigned from the ACBL.

Submitted respectfully,
Steve Moese
Secretary, District 11 Board of Directors.