THE GTOX BLUEPRINT

A SYSTEMATIC APPROACH TO GTO

POST-FLOP STUDY AND PLAY

V 2.0

CHAPTER 1: THE BUY-IN

Why All Players Can Benefit From Learning GTO[1]

A common perception in the poker community is that GTO should be learned and used solely by experts or high stakes players and that beginners and lower stakes players should only learn to play exploitatively/intuitively. In our opinion, this notion is misguided for a number of reasons.

GTO Defines Fundamentally Sound Play

GTO refers to a set of strategies between two players that maximize the average number of chips each player is likely to win against the other.[2] In other words, GTO strategies can't be improved upon to increase either player's expected win rate. As such, GTO tells us what the fundamentally "optimal" play is in a vacuum, which, in the long run, will tie against an optimal opponent and win against a suboptimal opponent.

There are few (if any) other games or activities where the conventional wisdom is that one should learn the fundamentals only *after* he has built habits, techniques and understandings based on something *other than* the fundamentals. When a player learns poker simply through trial-and-error and without a foundation built on the fundamentals, he runs the risk of developing poor habits due to the disconnect in poker between optimal *play* and optimal *outcomes*.

One of the most basic forms of learning is operant conditioning, where positive actions are rewarded (causing the positive action to increase) and negative actions are punished (causing the negative actions to be avoided). However, as a result of poker's chance elements, incomplete information and human unpredictability, positive actions (i.e. optimal plays) are often punished with a loss, and negative actions (i.e. blunders) are often rewarded with a win. This incongruence can lead a player to follow incorrect patterns and rationalize bad decisions unless he has a fundamental understanding of the

---

[1] We use "GTO" or "Game Theory Optimal" in the colloquial sense to describe a math and/or solver-based style of play that is primarily focused on developing unexploitable strategies against an opponent that is capable of adjusting its play.

[2] Presently known solving algorithms can only guarantee unexploitability for two-player zero-sum games. However, the current state-of-the-art algorithm that calculates nash-equilibrium strategies, counterfactual regret minimization, has also been successfully applied to non-two-player zero-sum games, in part, by utilizing the algorithm's function of identifying and eliminating strictly dominated strategies (i.e. strategies that always produce lower EV compared to an alternative strategy). *See e.g.,* R.Gibson, *Regret Minimization in Games and the Development of Champion Multiplayer Computer Poker-Playing Agents* (2014). *See also*, N. Brown and T. Sandholm, Science 10.1126/science.aay2400, 2 (2019).

game which informs whether the play was optimal in a vacuum independent of the outcome.

GTO is Effective No Matter the Stakes or Opponents

Since GTO is unexploitable, it is, by definition, not vulnerable against any player type, whether tight or aggressive, online or live, microstakes or nosebleeds. The most powerful bots and solvers, which use nash-equilibrium algorithms, are banned from use because they print money at all stakes. In fact, many believe that real time solver assistance and bots represent an existential threat to online poker because they are overwhelmingly effective at sucking profits out of the player pool.

Additionally, it is indisputable that the strongest players in the world tend to implement lines that are closer to the nash-equilibrium and that the weakest players in the world tend to implement lines that significantly deviate from nash-equilibrium strategies. Further, no aspiring poker player has an ultimate career goal of crushing the mean streets of $.01/$.02. So assuming that an understanding of GTO is needed to beat the highest stakes, a player that begins learning GTO early on, when the risk of loss is relatively low, will have an advantage relative to his peers that choose to hold off learning GTO until they reach higher stakes (which may never happen).

Exploitative Play Adds Complexity

Predominantly exploitative play, when effectively applied, is inherently abstract and it requires knowledge of domains outside of the mechanics of the game. In order for a player to know how to make an accurate exploitative adjustment, that player first needs to understand what the baseline play should be (*which will inherently be flawed if not based on GTO*). Then, the player must have a sufficient sample size of data on player or population deviations, which typically can only be effectively identified and internalized through experience over time. The player must then be able to accurately modulate the baseline strategy based on the player's understanding of these deviations, which can be difficult because exploitative considerations tend to be abstract and not objective or quantifiable. Finally, the player must be able to recognize when his opponent shifts strategies (or when he faces a new, different type of opponent) and, in response, the player must repeat these same steps to readjust. Given this dynamic, it can be very difficult, particularly for a newer player, to develop a baseline strategy if Hero is forced to continuously use a different type of strategy based on each unique opponent.

In a way, exploitative play could be viewed as an *advanced* strategy that should be reserved for players that first have obtained a firm grasp of basic fundamentals (i.e. some abstracted form of GTO). In contrast, since GTO strategies are calculated and are effective regardless of the opponent - GTO does not require any knowledge of strategic trends (which are constantly in flux), player types, reads, tells, human psychology or metagames, and therefore it decreases the quantum of variables that need to be considered and synthesized in the decision-making process.

Exploitative Players Can Benefit from GTO

While the notion that exploitative play can maximize profits over GTO is true in theory, there are many scenarios where having some knowledge of GTO strategies and how they are constructed can be beneficial to even the most exploitative players.

First, every player will encounter situations where they do not have a sufficient sample of hands on a particular Villain to accurately determine how the Villain plays in general or in specific scenarios. In those situations, defaulting to an unexploitable strategy until further data can be assessed will mitigate Hero's exposure to being outplayed, since any exploitative strategy is, by definition, exploitable itself. To put it another way, the primary shared goal of both a GTO-based and an exploitative-based playing style is to win as many chips as possible by executing EV maximizing strategies against the assumed strategies of Villain.

A GTO-based strategy assumes the opponent is playing GTO and an exploitative-based strategy assumes the opponent is deviating from GTO in a specific way. However, if these assumptions about how Villain plays turn out to be wrong, utilizing a GTO-based style should still win in the long run, whereas playing exploitatively will leave Hero open to being punished. As such, an adept player could maximize his EV by switching between a GTO-based and an exploitative-based style in different subgames depending on whether the player has sufficient data on Villain's tendencies in such subgames.[3]

Second, purely exploitative play doesn't actually tell us anything specific about how we should play our hands even when we are certain that our opponent is imbalanced. At some point, even the most aggressive exploitative strategies must include lines/hands that are played passively, and even the tightest exploitative strategies must include lines/hands that are played aggressively - and any system that identifies these lines/hands that is not based on fundamental principles of poker theory will be sub-optimal.

For example, assume that Hero is in the big blind in a single raised pot facing a Villain that over-bluffs. Villain triple barrels and shoves 80bb on the river. Hero is holding 4th pair. Should he call? What if Hero has Ace high instead? Queen high? Simply having knowledge that our opponent over-bluffs isn't sufficient to rationally decide how to play our particular hand. The reality is that even the most ardent exploitative players have

[3] For example, some Villains may play overly aggressive on all streets, but others may only play aggressively preflop and/or on the flop, but play much more passively on later streets. If Hero has a large sample size of earlier street play for a particular Villain, but limited data on later street play, Hero could play exploitatively on preflop and on the flop, but play a more balanced style on turns and rivers until more data is gathered. Additionally, if Hero identifies and exploits an imbalance in Villain's strategies, but Villain then adjusts, it may be advisable for Hero to switch to a "GTO-based style" until he can re-assess Villain's strategies.

some internal baseline they start with, and if that baseline is not derived from GTO it will be inherently flawed.

Furthermore, because a solver is simply an EV maximizing calculator, by studying solver solutions, we can gain insights into how optimal strategies that maximize EV are constructed based on whatever assumptions we input, even if those assumptions are not "GTO" themselves. In that regard, even a successful exploitative player can benefit from using solvers to obtain a fundamental understanding of the mechanics of maximizing EV, the relative value of similarly situated hands and how cutoff points between various actions (e.g. bet/check, raise/call/fold) should be shifted based on Villain's range/tendencies.[4]

Finally, although many view solvers as rigid and banal, solvers often use unconventional and unintuitive lines that can be cherrypicked and incorporated into an exploitative player's repertoire of plays. For example, to balance its range, the solver will find a variety of bluffs in almost any scenario - even in spots where most humans would instinctively think that no "natural" bluffs are available. The exploitative player could use the solver's low-frequency lines as inspiration to augment his arsenal of creative plays based on fundamentally sound, plus-EV strategies instead of arbitrary randomness. We refer to this as a "GTO-influenced" strategy.

[4] As a practical matter, the primary difference between GTO and exploitative play is simply where different cut-off or threshold points exist among various actions, whereas most of the generalized principles derived from poker theory that determine the relative value of hands remain applicable. For example, if Villain calls too loosely on the river, one of Hero's natural exploits will be to move the cutoff points of his bluffs to remove some weaker bluffs, and move the cutoff points of his value combos to add some thinner value bets. However, regardless of whether Villain calls at the correct GTO frequency or too loosely, the principles that determine which hands make better bluffing or value bet candidates remains the same (e.g. showdown value, card removal). As such, a player that is highly proficient at GTO will also likely be adept at exploits because he will have a strong sense of where cut-off points exist at the baseline, the most prevalent combos in his range, and the optimal way to shift cut-off the points along that range when facing an imbalanced Villain.

GTO can be Abstracted and Simplified

Possibly the most persuasive argument for the belief that beginners should not learn GTO is that historically, GTO has had a very steep learning curve that is difficult for novices to grasp. However, while some may think of GTO in a binary manner (i.e. either you are playing GTO or you are not), the more practical view is that GTO is a *spectrum* where alternative strategies are closer or further away from the nash-equilibrium. The reality is that even the most popular solvers on the market use abstractions and do not solve full games all the way to the nash-equilibrium, so their solutions cannot be considered 100% "true" GTO. In that regard, GTO can and should be abstracted to scale to accommodate the user's level of experience, and this is the primary function of GTOx.

By way of example, assume that in a 6max cash 100bb game, the CO 3bets the LJ and the flop comes Kh 4d 2s and the LJ checks. In this scenario, the solver's strategy, based on certain assumptions,[5] is for the CO to check 0.3% of the time, c-bet ¼ pot 70.31% of the time, c-bet ½ pot 22.72% of the time, c-bet ¾ pot 3.94% of the time and c-bet full pot 2.73% of the time. However, it is impossible for any human being to memorize frequencies at this level of specificity across all possible scenarios in no limit hold 'em. Accordingly, abstractions and simplifications are required, and the level of abstractions and simplifications we employ can be scaled based on the user's experience level. For example, an abstraction we can derive from this particular solution is that in a 3bet pot with moderate SPR, the in-position preflop aggressor should c-bet small with very high frequency on most dry king high boards. Although not as precise as the exact GTO solution, this heuristic is still consistent with GTO, and is something that can be learned and implemented by even a relatively inexperienced player.

[5] The exact strategy will change based on the inputs such as preflop ranges, rake and bet sizes.

The GTOx System

GTOx is a comprehensive suite of resources designed to assist players of all skill levels and playing styles with learning and implementing fundamentally sound poker strategies, consisting of four components:[6]

- The *GTOx Blueprint*, which provides a streamlined, systematic approach to studying and analyzing hands from a GTO perspective based on first order principles derived from years of rigorous solver work and data analysis.

- The GTOx Solver, which is a patent-pending cloud-based solver that calculates EV maximizing strategies for single hands and provides enhanced metrics and data visualizations to aid in learning core theoretical concepts.

- *GTOx Reports*, which provide aggregated analysis across multiple board textures that can be used to formulate generalized range strategies and attain a deeper understanding of the interaction of ranges with boards.

- *GTOx Solutions*, which are a library of solved flops in many different commonly encountered scenarios that may be explored using Simple Postflop, and played against using GTO Trainer.[7]

The GTOx Blueprint Philosophy: Systematic Simplification[8]

Due to the nearly infinite complexity of the game, all poker learning systems *must* incorporate simplifications and generalizations into their strategies. However, instead of simplifying the game through arbitrary methods, the GTOx Blueprint simplifies the game through a principled approach – by utilizing heuristics and abstractions while

---

[6] Additional modules to come.

[7] Users can change bet sizes and node lock turns and rivers for the GTOx Solutions using the free version of Simple Postflop. However, in order to change bet sizes or node lock flops, a full license is required, which is sold separately. Additionally, a separate license is required to use the GTOx Solutions with GTO Trainer. See www.simplepoker.com for more info.

[8] For the sake of brevity, this guide assumes the reader understands the basic mechanics, terminology and strategies of Texas no-limit hold 'em (e.g. hand rankings, betting mechanics, pot odds, etc.), which are freely available through various resources. Additionally, the GTOx Blueprint intentionally avoids discussion of math, which has been extensively covered in other poker theory books including *The Mathematics of Poker* by Bill Chen and Bill Chen and Jerrod Ankenman, *Expert Heads Up No Limit Hold'em Vol. 1* and *Vol. 2* by Will Tipton, *Applications of No-Limit Hold'em* and *No Limit Hold'em for Advanced Players* by Matthew Janda and *Modern Poker Theory* by Michael Acevedo, all of which are recommended for those that are seeking a deeper understanding of the foundations/proofs of GTO strategies.

retaining as much EV[9] from the GTO solution as possible. The GTOx system is not designed to completely revamp a player's playing style so that he plays precisely in accordance with the solver, which is impossible. Rather, the primary focus of the GTOx system is to identify universal principles of EV maximization and highlight those areas where a player has the biggest leaks, providing insight into why such leaks exist and offering principles that the player can use as touchstones to mitigate such leaks.

As a caveat, it is important to note that any system of simplification, including the GTOx Blueprint, will, by definition, deviate from GTO and therefore will be inherently "flawed" and result in a certain degree of "errors". In fact, all human developed GTO methodologies and systems are, by their very nature, based on a fiction since solvers generally do not use defined strategic tactics. However, absent a better alternative, as the saying goes, we shouldn't allow the *perfect* to be the enemy of the *good*. The reality is that even the most powerful supercomputers utilize abstractions when calculating GTO strategies. For example, Pluribus (the most advanced multi-way poker AI published to date) has a pre-defined, coarse-grained "blueprint" strategy which is used as a starting point and is then refined during actual play. This simplification is necessary because it is impossible for any computer (or human) to store the strategies for the entire game of poker in memory.

In essence, the GTOx Blueprint is a blueprint strategy for humans. It imputes a humanistic rational to patterns derived from solver solutions so that the user can utilize a principle-based approach to *approximate* GTO rather than relying on rote memorization of sims or complex calculations. Most of the principles and concepts presented in this guide are not new- the innovation the GTOx Blueprint provides is organizing these principles and concepts in a streamlined, rational way that can be applied in a systematic manner, even by lesser experienced players and across all types of no-limit hold'em games.

The GTOx Blueprint does not provide advice on how to play specific hands – a solver already does that with nearly perfect precision. The primary intent of the GTOx Blueprint is to provide a practical guide which players of all levels and styles can use to interpret and understand solver solutions - with the ultimate goal of, not *replicating* GTO, but rather, *understanding* poker theory in an intuitive way so that it can be used as one of the tools in your toolbox to improve your own unique game.

Lastly, since the GTOx Blueprint is a simplified system of analysis, it does not address many of the subtleties and nuances that are ubiquitous in GTO, which are extremely *context* dependent and very sensitive to input. In order to truly master GTO, you will

---

[9] EV stands for "Expected Value", which is the average number of chips a player is expected to win by the end of a game (i.e. at showdown or when one player folds) given each player's strategies. Since the primary goal in poker is to win as many chips as possible, our goal is to take actions that maximize EV.

need to do a lot of legwork, consisting of solver study, to fill in the contours and to train your intuitive understanding of where specific strategic lines are drawn. Ultimately, the *solver* should be your coach – no human will ever be more accurate than a solver when it comes to identifying optimal strategies. To draw an analogy, a human trying to teach the *optimal* strategy for a specific poker scenario based on his own experience or intuition is like someone trying to tell the time of day by looking at the sky and measuring the position of the sun relative to the horizon. That person may be able to give you a close approximation of the actual time - but why even bother if you have a fine-tuned quartz watch on your wrist? The solver is that watch.[10]

[10] This is not to say that coaches are not valuable for players seeking to increase their win-rate. There are many other aspects to winning poker that go beyond understanding fundamental theory, such as exploitative adjustments, bankroll management, allocation of study time, mindset, game selection, player pool tendencies, tournament meta-strategy, etc. However, these aspects are outside of the scope of this document, and when it comes to identifying strategies that maximize EV given set assumptions, there is no human that can approach the precision of a solver.

CHAPTER 2: OVERVIEW

Fig. 1. The GTOx Solver. The GTOx Solver presents outputs in an interactive data-visualization dashboard which can be used to analyze the solver's outputs through a system of cross-filtering across different charts. The information in the dashboard is curated, structured and enhanced to assist the user in understanding GTO strategies.

The Tri-Level System

The GTOx Blueprint is structured in three sequential levels: [11]

Level	Focus
MacroAnalysis	The overall tendency of Hero's range.
MesoAnalysis	The overall tendency of Hero's hand class.
MicroAnalysis	The probability altering characteristics of Hero's specific hand.

[11] Although for the sake of simplicity, the GTOx Blueprint breaks down the hierarchy of granularity into 3 levels (Range>Hand Class>Specific Combo), throughout this guide, references will be made to two additional, intermediary levels, referred to as "Hand Segments" (i.e. full houses, flushes, sets, two pairs, top pairs, under-pairs, ace highs, air, draws) and "Combo Clusters" (i.e. AKs, KJo, TT, etc.) since the solver will often balance its range not just across the range and hand classes, but within individual segments and clusters as well. When segments and clusters are included in discussion, the hierarchy of granularity is as follows: Range>Hand Class>Hand Segment>Combo Cluster>Specific Hand.

This structure serves four core purposes: (1) Prioritization, (2) Simplification, (3) Compartmentalization and (4) Unpredictability/Flexibility.

Prioritization

The order of the GTOx Blueprint analysis is important because in GTO, the tendency of Hero's range (*the MacroAnalysis*) will often supersede the tendency of Hero's hand class (*the MesoAnalysis*), and the tendency of Hero's hand class (*the MesoAnalysis*) will often supersede the specific characteristics of Hero's hand (*the MicroAnalysis*). As such, utilizing separate levels of analysis, which force Hero to consider the strategy of his range before considering the strategy of his specific hand, helps the user understand how GTO strategies are constructed, where the same hand may be played completely differently depending on the context.

Simplification

The tri-level system is ordered in increasing levels of granularity, and to simplify strategies, the GTOx Blueprint provides users with the flexibility to make decisions at higher levels of abstraction when doing so will not result in significant EV loss. This approach greatly simplifies the game in a rational way as there are many scenarios in poker where nearly the entire range or an entire class can be played in a similar manner without significant (or any) EV loss – and therefore drilling down to the most granular level may be counterproductive since the solver will often make hair-splitting decisions at the MicroAnalysis level based on how the ranges specifically matchup combo-by-combo, which is impossible to be accurately implemented by humans.

Compartmentalization

Given that there are a multitude of factors and principles that drive GTO strategies, utilizing a system that organizes the decision-making process is critical. The GTOx Blueprint creates a logical structure by breaking-up these factors into three separate compartments of analysis, which allows Hero to focus on the specific factors that relate to the applicable level of analysis instead of attempting to synthesize all of the relevant factors at once which can become unwieldly.

Unpredictability/Flexibility

Instead of requiring memorization of specific "moves" which can be identified and countered by opponents, the GTOx Blueprint provides for a principled approach where each level of analysis informs the next level, leading to a final decision that is a systematic aggregation of all of the most important relevant factors. This makes Hero's play inherently unpredictable because similar spots may be played differently based upon the cumulation of factors from all three levels, some of which are unknown to Villain. Additionally, since the GTOx Blueprint is not rigidly built on formulaic "if *x,* then *y*" conditions, it allows for flexibility and can blended with exploitative play (while

still being tethered to GTO fundamentals) by allowing the user to bias his strategy towards aggression or conservation when doing so will not result in significant EV loss.

EV Regret

EV is the most important metric in poker because it tells us how many chips on average a player will win by taking a particular action against a perfect opponent. Because of its importance, every competent player that has ever used a solver, avails to EV as a measure of the degree of his errors. "Punts" occur when a player takes an action that results in a significant EV loss compared to an alternative action. A "minor error" occurs when a player takes an action, which may not be the optimal one, but results in substantially the same EV as the best action.

Given the complexity of poker, it is impossible for any player to correct all minor errors in his game, which may be the result of minutia relating to how the ranges specifically match up, so devoting significant time and attention on these details typically bears little fruit. On the other hand, users can significantly improve their game by reducing or eliminating punts, so spending time to analyze these spots can result in major EV gains.

However, historically, players could only measure EV loss on a combo-by-combo basis. GTOx is the first system that allows players to measure max EV loss across groups of combos as defined by the user through the invention of "EV Regret".[12] EV Regret measures the EV loss between any action and the best alternative action (if any), not just for one combo, but across all combos in a group of hands selected by the user, and then returns the maximum EV loss across all such combos, divided by the size of the pot.

[12] EV Regret is analogous to the concept of "regret" in game theory, which measures the amount by which an alternative strategy outperforms a player's current strategy. Nash-equilibrium strategies can be discovered through "regret matching" whereby the player gradually alters his strategy to skew his probability distribution proportionally towards actions that have the most positive regret. This process continues until the average regrets of all players equal zero, meaning that further adjustments will not increase EV, which occurs at the nash-equilibrium.

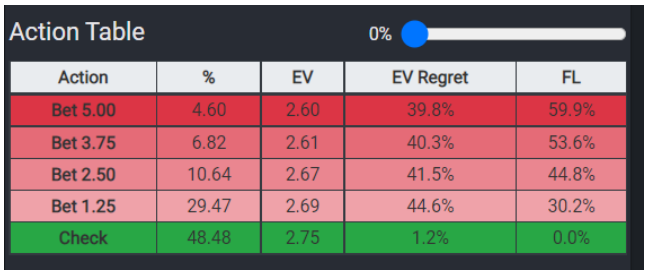

Fig. 2. EV Regret. EV Regret calculates the maximum amount of EV (divided by the pot size) loss for all combos in the player's class compared to the GTO strategy. The user can create his own "class" of hands by utilizing the cross filters in the dashboard (e.g. combinations of hand segments, draws, equity tiers, suits, blocker/unblocker scores, etc.). In the above example, if the player bets 5.00 with all of the combos in the selected class, then there is at least one combo that loses 39.8% of the pot compared to the optimal action for such combo.

Measuring EV loss, not just for individual combos, but across multiple combos in user-defined classes is useful, because no player devises a unique strategy for each individual combination in his range. Instead, players devise strategies for groups or "buckets" of hands (e.g. top pair, strong kicker, flush draw, nut straight, trash, etc.). Accordingly, having the ability to quickly test the max EV loss across all of the combos within such groups or "buckets" allows the user to not only test the validity of his strategy for his specific hand, but across all similar hands that he would play the same way.

For example, assume that a beginner devises a strategy to c-bet range 1/2pot on all Ace High boards. Although certain individual combos may prefer to c-bet 1/2pot on Ah4h5h in a single raised pot Button versus Big Blind, through the use of EV Regret, the user can quickly identify the fact that at least one combination in the range loses significant EV using this simplified strategy. This signals to the player that his strategy likely needs refinement.

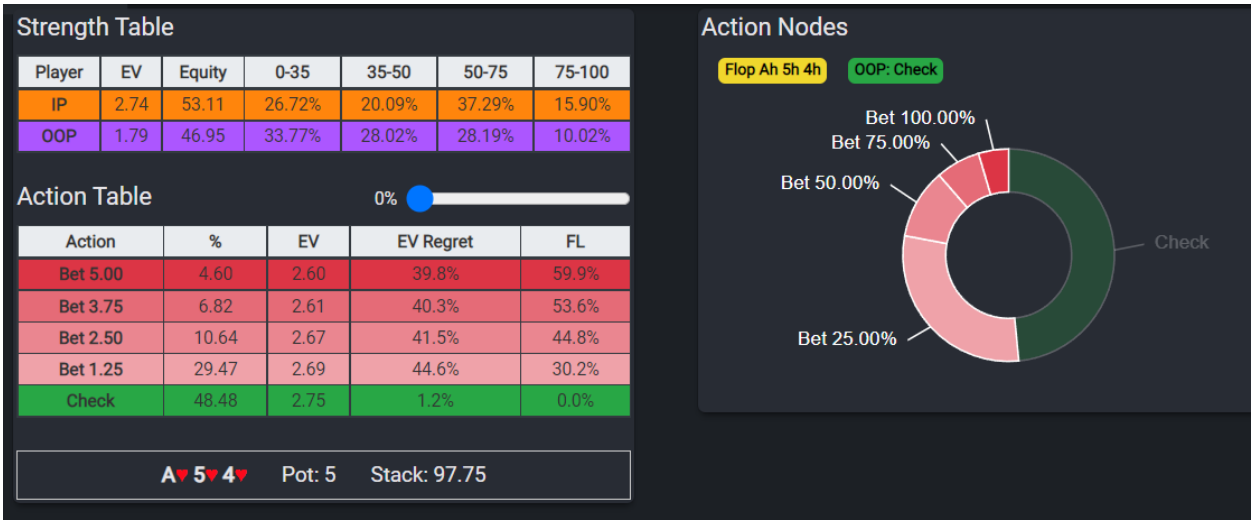

Fig. 3. EV Regret Example. In the above example, although 10% of the Button's range prefers using a ½ pot bet on this Ah5h4h board, there is at least one combo within the Button's range that loses a massive 41.5% of the pot betting 1/2pot compared to an alternative action. Accordingly, it would make sense for the Button to refine his c-betting heuristic by breaking down his strategy more granularly. The pie chart on the right shows the average actions taken by the range as a whole. The highlighted action can be clicked on to navigate to the next node.

Because EV Regret measures EV loss across user defined buckets of hands, it can be used to quickly identify whether a player's strategic logic contains any major errors.[13] As such, EV Regret can also be used to help simplify GTO strategies by identifying whether abstracted strategies lose significant EV compared to the GTO solution.

Due to poker's complexity, using abstracted strategies are a must. Accordingly, to acclimate one's learning of GTO, we recommend that users devise strategies at the highest levels of abstraction possible,[14] only proceeding to the more fine-grained levels if doing so will win non-"minimal" additional EV in accordance with the following:

EV Regret[15]	Recommendation
Only one action has minimal EV Regret	Follow that action with the entire range or class[16]
Multiple actions have minimal EV Regret[17]	For a more simplified approach, (a) mix between such actions using a randomizer weighted towards the action with the highest frequency or (b) use exploitative considerations
	For a more refined approach, drill down to the next level of analysis
No actions have minimal EV Regret	Drill down to the next level of analysis

[13] Note that EV Regret does NOT measure the exploitability of a player's strategic logic, which requires an assessment of frequencies. *See "A Note on Frequencies" below.*

[14] The most abstracted strategies are the simplest to remember and implement. For example, cbetting range is a common abstraction that is used by many because it allows a player to use one simple strategy across all combos in his range. However, the more abstracted a strategy is, the more exploitable it is and the more EV it will lose from the GTO solution. As a player progresses, he should attempt to make his strategy more nuanced and refined, devising unique strategies for smaller and smaller buckets of hands. This makes a player's strategy more disguised, balanced and unpredictable.

[15] Note that in certain circumstances when a hand has very low weight (e.g. very low probability of being within the GTO range), the solver may take an action with high frequency that has significantly negative EV. In such circumstances, it is generally recommended to assume that the combo and play is not part of the GTO game tree. The user can move the Wt. Decimal slider to eliminate hands with low weight (*see Fig. 4*).

[16] In certain circumstances, the solver may show only one action with minimal EV Regret, but that action is played with a minority average frequency. This may occur when most of the hands within the same class or range are being played with a mixed strategy across actions with identical EVs, but one or a few hands within the class or range strongly prefer one of the actions over the others, thus resulting in the other actions indicating a high EV Regret for the class/range as a whole. In those cases, although taking the action with minimal EV Regret with 100% frequency across all combos in the range or class should not lose in expectation against a GTO opponent, such a strategy may be exploitable. Accordingly, in these circumstances it is generally advisable for more advanced users to drill down to the next level of analysis.

[17] Given that solvers generally cannot solve full game trees all the way to the nash-equilibrium, there will be situations where the solution for a particular hand will show mixed frequencies for alternative strategies with different EVs. Technically, this is an error since at the nash-equilibrium, the optimal strategy will always be the strategy with the highest EV and strategies are only mixed when the EVs for two or more actions are identical. When these spots arise in GTOx, as a general rule of thumb, it is recommended that users assume (a) the EVs for strategies that show significant mixing of frequencies to be equal, or (b) if one action has disproportionately greater frequency and higher EV than all other options, the nash-equilibrium strategy is 100% frequency for that action.

Whether EV Regret is "minimal" depends entirely on each individual's experience and skill level. Since EV Regret is standardized as a percentage of the pot,[18] each user should set his own EV Regret threshold. As the user becomes more experienced with GTO, the goal should be to continually lower the EV Regret threshold to capture additional EV from the GTO solution. In this way, the GTOx Blueprint is scalable from beginners to experts and focuses on correcting the user's biggest leaks first and then adapts to each user's unique learning progression.

A Note on Frequencies

Ultimately, the GTOx Blueprint makes a deliberate tradeoff - choosing systematic simplification over optimal mixing of frequencies. We believe that simplifying the game with a focus on retaining EV but compromising on precise GTO frequencies is rational for several reasons.

For one, GTO's mixing strategies are incredibly complex and often times unintuitive. Accordingly, in order for a user to approach mastering optimal frequencies, he will need to run and memorize many different types of simulations, which, given that there are nearly an infinite number of different possible scenarios in poker, is not feasible.

Secondly, since GTO always takes the action with the highest EV, even if by a tiny fraction of a big blind, GTO frequencies are highly sensitive to different range and betting inputs. Slight changes to these variables will often result in dramatic differences in the mix of strategies, particularly as you traverse down narrow branches of the game tree.[19] For example, the solver may play the exact same combination in the exact same spot differently if the combination has a high weight/frequency versus a low weight/frequency, due to the solver's propensity to balance its range. However, in reality, unlike the solver, we do not know our opponent's (or our own) exact ranges or strategies by weight, which means the actual GTO mixed frequencies for a real-time scenario will never be knowable, and, furthermore, the ability of a human opponent to exploit less than optimal frequencies is extremely limited, so a dogmatic adherence to GTO frequencies should be avoided.

---

[18] The GTOx system standardizes EV Regret by the size of the pot instead of big blinds because an EV loss in a large pot is generally viewed as a smaller mistake compared to the same EV loss in a smaller pot. Another possible way to standardize EV Regret is by dividing the EV loss by [*pot + bet size*]. However, when EVs are close, including the bet size in the denominator may skew EV Regret scores towards larger bet sizes in a manner that may not be aligned with the solver's allocation of frequencies.

[19] Conversely, EVs tend to be more robust when simulation inputs are changed. For example, if moderate changes are made to the preflop ranges input in a simulation, many river bluffs may continue to have 0 EV or near 0 EV, however, the frequencies of specific bluffs may change dramatically due to blocker effects.

Third, since solvers use abstractions[20] and do not calculate all the way to the nash-equilibrium, we cannot be certain that a solution's frequencies are absolutely "correct" and they are in fact sometimes "wrong" (e.g. the solutions will at times take actions that do not maximize EV and solutions from different solvers will sometimes show different frequencies in the exact same spots).

Fourth, solvers do not recognize and cannot independently compute the effect of human imbalances and weaknesses, which can be incredibly important. Rather, solvers calculate GTO strategies assuming each player is playing a perfectly unexploitable strategy based on the parameters of the simulation. As such, solvers will often allocate frequencies to certain borderline hands primarily based on small differences in probabilities (e.g. blockers/unblockers/backdoors) which will result in hair-splitting decisions between different combinations that have very similar EV and equity – and this delicate balance will usually be completely destroyed when one of the player's strategies shift from GTO. In that regard, it would be imprudent to completely ignore reliable exploitative considerations when present in favor of a pure GTO-based approach created in a vacuum, particularly in close decisions. To maximize "true" EV, players should utilize all pertinent information available to them, which the solver is unable to do.

Finally, we should keep in mind that the most important metric in poker is EV and that solvers determine strategy by choosing the actions that are expected to maximize profit over the long run. That is, between two actions that have different EVs, GTO will *always*[21] choose the action with greater EV. In other words, GTO frequencies *always* follow EV, with the one exception being where the EVs between two or more actions are identical, in which case, the frequencies will generally be mixed based on how the players' ranges matchup to enforce the indifference principle.[22] In that regard, if, for the

---

[20] Discovering the "true" GTO strategies for a particular hand scenario would mean that each player, from preflop through the river, has the option to choose every action and bet size available within the rules. However, this would make the game tree so large that it would be computationally impossible to solve based on current technology. So instead, when setting up a simulation, we use abstractions by assuming certain parameters of the game, such as preflop ranges of the players and discrete betting sizes, that are designed to approximate the full game. The solver is able to solve this abstracted version of the full game to an accuracy that approaches the nash-equilibrium and we can then analyze these solutions to understand the principles that are driving optimal strategies based on the assumptions we have chosen.

[21] Note that in certain circumstances, solvers may not in practice allocate frequencies to strategies that maximize EV, but this is merely a consequence of the solution not being solved all the way to the nash-equilibrium.

[22] Generally, the solver will mix strategies when there is no single action available that will maximize EV when taken with 100% frequency. For example, assume Hero has two options available at a given decision point – Action A and Action B. If Hero takes Action A 100% of the time, Villain will be able to make an adjustment which will increase the average EV of Action B over Action A. However, if Hero then counter-adjusts by taking Action B 100% of the time, Villain can then make another counter-adjustment which will increase the average EV of Action A over Action B. In this scenario, the solver will mix Hero's frequencies

sake of simplicity, a player that is learning GTO focuses primarily on maximizing EV gain/minimizing EV loss, he should directionally converge towards optimal frequencies as he gradually reduces his EV Regret threshold over time.[23] In essence, this is similar to the process that solvers use to identify GTO strategies. Most commonly used solvers learn through simple trial and error – first attempting every possible line with equal probability and then gradually shifting its strategies towards the lines that maximize EV and shifting its strategies away from the lines that lose EV until it reaches the point that further adjustments will not increase EV against its opponent - and it is at that point that the nash-equilibrium is reached.

In light of the foregoing, spending time to memorize exact GTO frequencies can be counterproductive and, instead, this time can be allocated more efficiently to other areas. Given poker's enormous complexity, we recommend that GTO novices view overall strategies in a binary manner – *pure* or *mixed*, with the initial focus on identifying the pure strategy spots, as these are the areas where people tend to lose the most EV. If Hero is able to identify most pure strategy scenarios, by default, the remaining spots will be mostly mixed, and it generally will be more difficult for Hero to make a massive EV error in these mixed spots even if he is unsure of the correct strategy and many opponents will not be capable of taking advantage of Hero using less than optimal frequencies.[24]

This is not to say that understanding mixed strategies is unimportant.[25] The EVs generated by the solver assume each player's actions are taken in accordance with the frequencies of the solution, and if they are not, the actual frequencies used may be exploited and therefore lose in expectation compared to the precise GTO strategy. For example, in certain situations, Hero's river bluffs may all have nearly 0 EV, but if Hero

---

between Action A and Action B to maximize average EV overall, and the exact mix makes Villain indifferent from an EV perspective to two or more actions. *See footnote 27 for additional discussion of the concept of indifference.*

[23] Although, technically, simply minimizing EV loss for each action will not lead a player to the correct frequencies for mixed strategy spots and therefore may leave the player's overall strategy exposed to exploit over the long run, as a practical matter, it will be very difficult for a human opponent to identify and take advantage of such frequency imbalances as long as the player uses some type of mixed strategy for these spots. It should also be noted that using a pure strategy for a range or class will make the solution for subsequent streets less accurate since the strategies for subsequent streets assume that the player utilized GTO frequencies on prior streets.

[24] Technically, mixed strategies should have the same EV at the nash-equilibrium. However, since solvers are not able to solve the entire game tree all the way to the nash-equilibrium, there will often be spots where the solver mixes it play between actions that have differing EVs. In these situations, it may be the case that one or more of such actions actually will have 0% frequency at the nash-equilibrium. However, taking actions with 0% frequency but comparable EV to the action taken with 100% frequency, will generally result in smaller leaks compared to taking other actions with 0% frequency but much lower EV.

[25] Focusing on learning precise frequencies can be thought of as the final step in "mastering" GTO after a player no longer, or very infrequently, makes large EV errors.

decides not to bluff at all, his value combos will lose EV against an aware Villain that reduces or even eliminates his bluff-catching. As such, if you play regularly against the same capable opponents, keeping frequencies somewhat in check will also be important since it is possible to play with minimal EV loss but still be very exploitable if your frequencies are imbalanced. However, for most players, instead of focusing on frequencies of individual solutions, focusing on aggregated frequencies (e.g. overall bet/call/fold frequencies on the flop, turn and river for single-raised, 3b and 4b pots) will be more important and this can be tracked by using a HUD or a trainer that keeps statistics on overall aggression. Additionally, particularly for more advanced players, when constructing strategies, we recommend keeping in mind the general principle of balance – that is, ensuring that hands of varying strength are retained in various branches of the game tree.

What We Can Learn From Solvers

As a final general note before we dive into the nuts and bolts of the GTOx system, it may be helpful, particularly for those that may not have much experience with GTO, to take a step back and briefly touch upon a few overarching concepts to keep in the back of our minds by highlighting what we are primarily aiming to accomplish through solver study.

Thinking in Ranges

In a vacuum, the fundamental goal in each and every hand is to maximize the number of chips a player has by the end of the hand.[26] To achieve this goal, one of the core analytical tasks a player must undertake is to estimate how far ahead or behind such player is from his opponent. The further ahead a player is, the more chips that player generally will want to be placed into the pot. The further behind a player is, the more inclined that player should be to give up or bluff (depending on how far ahead his opponent is). This, of course, raises the question of how we determine who is ahead or behind.

Given that we do not know Villain's actual cards, in order to estimate where we stand relative to our opponent, we must assess the likelihood of Villain holding certain combinations, or, in other words, we need to define Villain's "range". Simply put, a player's "range" consists of all of the likely hands the player may be holding at a given decision point. A player's range is shaped based upon his prior actions, starting with preflop and as further delineated by each of his subsequent actions.

As such, *before each decision*, it is imperative for Hero to consider (a) how Villain's prior actions up to that point define Villain's range by discarding or discounting combos from Villain's range that were unlikely to have taken Villain's actions and (b) how

[26] In many instances this will be fewer chips than the player started with, which is why folding is often a viable option since it cuts off potential future losses.

Villain's range is expected to respond to Hero's actions by putting himself in Villain's shoes.

However, we cannot simply focus on the composition of our *opponent's* range. Instead, we must also take into account Villain's likely perception of *our* range as well since Villain should also be engaging in the same exercise of trying to determine our range. In an abstract sense, poker can be thought of as a game where players place wagers based on the *perceived probability*[27] that they are holding the best hand, and the more accurately a player is able to define both players' ranges, the more precisely that player will be able to calculate this probability. Accordingly, *before each decision*, it is also imperative for Hero to consider how his own prior actions define his range, as Hero's range should influence Villain's actions.

And this is one of the reasons why GTO is often viewed synonymously with the concept of *balance* – that is, in our practical implementation of GTO, we want to make it as difficult as possible for Villain to identify our cards and strategies, and this is largely achieved by playing many different holdings - strong, medium and weak - in a similar manner. This helps keep a variety of segments of our range in play, thus providing more cover for our actual hand on different runouts.[28] Simply put, if we fail to balance our range, our range becomes more definable by discerning opponents, thereby making us

---

[27] We refer to the *perceived probability* in lieu of the *actual probability*, because each player's cards are hidden, and therefore a player can create the *perception* that his opponent's *probability* to win is low (e.g. by leveraging range advantage and bluffing).

[28] In addition to the benefit of disguisal, or information hiding, a perfectly balanced strategy is unexploitable in that it prevents Villain from increasing his EV by bluntly choosing one strategy over another with certain *borderline hands* – in other words, a perfectly balanced strategy forces Villain to be *indifferent* between such actions with such borderline hands because if Villain is not indifferent, then it means he can shift to a pure strategy and increase his expected payoff, thereby reducing Hero's payoff and violating the definition of the nash-equilibrium. For example, when Hero is facing a decision to bet or check, if Hero uses an optimally balanced bluffing/value betting strategy, Villain will not be able to exploit Hero through overcalling or overfolding and instead, to avoid being exploited himself, Villain will be forced to "top up" his defending frequency with a precise proportion of marginal bluff-catchers that are indifferent between calling and folding. Along the same lines, when Hero is facing a decision to defend or fold, if Hero uses an optimally balanced defending/folding strategy, Villain will be unable to exploit Hero through over or under bluffing and instead, to avoid being exploited himself, Villain will be forced to "top up" his bluffing frequency with a precise allocation of marginal weaker combos that are indifferent between bluffing and checking. That being said, although this indifference principle is central to GTO in theory, in terms of a practical application, most readers would benefit from simply keeping in mind that in order to avoid having major leaks that can be exploited by opponents, you should generally try to play with some degree of balance, such as by pairing strong hands with weak hands in various action frequencies, by playing some weak hands aggressively as bluffs and others more passively, calling with some marginal bluff-catchers and folding others, and slow-playing with some strong hands when your overall range is doing a significant amount of checking, etc. The allocation of these ratios should be guided primarily by the size of the bet in relation to the pot and the composition of the player's ranges. Studying many different solver solutions will aid the user in obtaining an intuitive understanding of where the optimal lines and points of indifference exist based on these variables.

easier to play against and exploit. In this regard, philosophically, one practical goal of a real-life implementation of GTO can be thought of as attempting to define Villain's range as narrowly as possible, while at the same time, obfuscating Hero's own cards (and strategies) by balancing his range. GTO defines what the optimal balance is in this regard and GTOx can be used as a key to understand how to *approximate*[29] this balance.

Fig. 4. Ranges. The ranges of the Out of Position Player (OOP) and In Position Player (IP) are depicted by these heatmaps, with the shade of each cell denoting the relative frequency the combinations within such cell are present in the GTO range. The user can view the GTO strategies for the player to act by clicking on the "Actions" button or by hovering over each shaded cell in the range. The ranges of each player will narrow at each decision point based upon each player's respective actions. Given that range construction is perhaps the most critical skill to master in poker, focusing on building an intuitive understanding of the shape of each player's range in different spots should be a priority. The user can filter out lower weighted hands by using the "Weight Decimals" slider. The default setting is "0", where no filter is applied. The user can filter the dashboard to only include hands with a value within "1" decimal place up to hands with a value within "8" decimal places. This slider is particularly useful on later streets, where hands with very low weight may take non-GTO actions. Clicking on the "Revert" button allows the user to deselect combos from the range to create refined buckets.

Principles and Mechanics of EV Maximization

One of the most common criticisms of GTO is that no human player actually plays nash-equilibrium strategies and therefore the assumptions that underlie solver simulations are not realistic and can lead to play that does not maximize profits. To a certain degree, there is merit to this argument. For example, on the river, the solver may call 100% of its AdKs combos and fold 100% of its AhKc combos in a scenario facing a large overbet, despite the fact that both hands have identical showdown value, and this may be the result of certain suits blocking or unblocking Villain's betting range on the river, which itself was a result of certain suits blocking or unblocking holdings that Hero had in his

[29] We emphasize that our goal should be to *approximate* the solver's balance because the level of balance achieved by the solver is impossible for humans to replicate, and is of marginal value against human opponents, who do not know our precise ranges or strategies, are unable to accurately track our frequencies over the long run in every situation and can't quantify or calculate odds on the fly with the same precision as a computer.

range on the turn and so on and so forth - tracing all the way back to the initial actions taken preflop. However, it is extremely unlikely that a real-life opponent will ever utilize suits with anywhere near the same precision as a solver, and therefore assuming that Villain's range is identical to the solver's range in a given spot can result in a significant opportunity cost.

Does this mean that the GTO strategy is useless in this context? Well, no. For one, the closer a player is able to implement a GTO strategy, the less likely the player will be vulnerable to exploit regardless of the opponent since a player approximating GTO does not need for his assumptions about his opponent's strategies to be accurate in order to for his strategies to be effective.

Furthermore, beyond emulating GTO strategies to avoid being exploited, studying simulations can also serve the purpose of illuminating how EV is generated. That is, although it may not be very useful to memorize *what* the solver's strategy looks like in a specific scenario, it can be extremely useful to understand *why* the solver used its particular strategy/counterstrategy in that scenario. In other words, it may not be productive to memorize that the solver is calling with AdKs and folding with AhKc in our specific example from above, because it is very unlikely we will ever encounter that exact same spot again. However, it would be useful to recognize the fact that the reason *why* the solver chose to call AdKs and fold AhKc was because (a) the caller needed additional hands to defend given the ranges, pot size and bet to prevent Villain from printing money by over-bluffing, (b) AK beats most of the Villain's bluffs and (c) AdKs *blocks*, and AhKc *unblocks*, many of the value combos in Villain's range.

We should keep in mind that solvers do not know what "GTO" is. Although technically, the nash-equilibrium strategies as a whole are break-even assuming that all players are playing perfectly and switching positions, the true value of GTO lies in the fact that these break-even strategies are achieved by playing each combo within each player's range in a manner that maximizes profits. In other words, solvers calculate the strategies that maximize EV based on whatever assumptions we use as inputs – whether those assumptions are based on GTO or anything else.

In that regard, solvers are useful because they can help us identify the *mechanics* of how EV is maximized, including the *principles* that commonly drive how optimal strategies and counterstrategies are constructed across different scenarios based on certain assumptions – and these *mechanics* and *principles* can be applied in our own games by adjusting those assumptions in light of our actual opponents and how they play their ranges (e.g. instead of calling with AKo, we can choose to call with some other marginal combination which blocks the assumed value range of our actual opponent, whatever that may be).

And while the systematic, stepwise structure of the GTOx Blueprint, which has been designed to focus users on these mechanics and principles, may seem overly rigid or complex to use in real time to some, we need to keep in mind that playing optimal poker

is a skill, and just like any other skill, there are basic building blocks which must be learned before a player is able to successfully implement that skill.

To use an analogy, a seasoned golf instructor wouldn't teach a beginner how to swing by simply handing him a club and telling him to go hit some balls. Although it is possible for the beginner to eventually learn how to swing properly through trial and error, there is a more efficient path. To teach a beginner to swing, a good golf coach will show the beginner, step-by-step, each important fundamental building block of a solid swing, such as how to grip the club, where the feet should be placed, how the arms and hands should be positioned in the backswing, how to use the hips, the ideal finishing position, etc. etc. When the beginner is first learning, he will need to deliberately think through each of these steps, but as he gets more and more repetitions under his belt, eventually some or all of the steps will be internalized to the point where the player will no longer need to consciously think of each step and instead, he'll be able to just line up and swing.

Leak Identification

In addition to the function of identifying optimal strategies, solvers can also provide us with an objective way to measure our biggest leaks through the metric of EV. As noted above, there are many reasons why prioritizing EV over mastery of precise GTO frequencies may be warranted for many different types of players, especially for those that have historically been excluded from the GTO discussion such as beginners and exploitative/intuitive based players.

For beginners, focusing on EV over precise frequencies allows for highly abstracted strategies, giving novices the space to get their feet wet in the realm of solvers from a 500 foot level before digging into all of the complicated nuance and minutia of GTO that experts use to gain slight edges over their opponents. Specifically, beginners can utilize solvers to objectively identify spots where they are losing the greatest amounts of EV from GTO strategies and then make corrections by "pruning" those branches of the game tree. And as the beginner whittles away at these spots and gradually reduces the incidence of large EV errors from his game, he will, as a natural byproduct, eventually get directionally closer and closer to following the "correct" GTO frequencies, since GTO frequencies "always" follow EV.

For exploitative players who may not have an interest in playing a perfectly balanced style and would rather take maximum advantage of their opponents' imbalances, the EVs derived from GTO solutions can also be used as a constraint, or outer boundary, on the degree to which they strategically deviate from the nash-equilibrium. This outer boundary can ameliorate some of the dangers of exploitative play's dreaded slippery slope. The notion of "playing the player and not the cards" is becoming outdated as more players at the top of the food chain are utilizing GTO strategies, which in turn causes balanced playing styles to trickle down to lower and lower stakes. Opportunities for exploit are still available at all levels, but having a quantifiable method to tether an

exploitative player's play, even loosely, to fundamentally sound strategy will reduce that player's vulnerability to being exploited himself.

Given that GTO strategies are usually heavily mixed, there are typically many different lines within the confines of GTO frequencies that can be biased towards aggression or conservation without major EV loss. In that regard, even players who play maximally exploitatively can utilize solvers to identify exploitative lines that are within the proximity of optimal strategies and exploitative lines that are losing significant EV from the GTO solution, and eliminate, or reduce the frequency of, the latter, while still playing a predominantly exploitative strategy. For example, when Hero has an exploitative read on a particular opponent or player pool, one "quick" exploit Hero could easily implement would be to simply shift all of his strategies that would otherwise be mixed/indifferent at the equilibrium to pure strategies. This would allow Hero to take advantage of his opponents' imbalances while at the same time retaining a strategy that will remain effective against a balanced, non-exploiting opponent.

Using the AKo bluff-catching example from above, although the solver may choose to call AdKs at a high frequency and fold AhKc with a high frequency when facing a large overbet on the river, often times the EV difference between the options will actually be very close – perhaps a small fraction of a big blind. If Hero knows that his player pool generally under-bluffs the river, Hero could make an exploitative play by folding *all* of his AKo combos regardless of suit, as well as all of his other bluff-catchers that are indifferent between calling and folding at the equilibrium. However, Hero will likely still need to choose other holdings that aren't the stone cold nuts to call with to mitigate his vulnerability to bluffs - and the GTO solution can be used as a guide to dictate where Hero draws this line by identifying the combos within his range that lose *significant* EV when folding. By focusing on EV in this manner, the exploitative player can continue to utilize exploitative data to influence his decisions, while at the same time being kept in check by an objective outer boundary to reduce his own exposure to being exploited.