“Got to be In It to Win It”:

Impact of Individual Risk Aversion on Ability to Monetize Profitable Investment Opportunities[1]

By Vladimir Ragulin,
Independent[2]

Abstract:

This research note is inspired by a simple but illuminating behavioural finance experiment by Victor Haghani and Richard Dewey[3], built around flipping for profit an unfair coin with known heads/tails probabilities. Using this elementary framework with fully known game parameters and well-understood win/loss probabilities the paper shows that our ability to take advantage of this clearly profitable opportunity greatly depends on our attitude to risk, our outside financial resources as well as on financial strength of “the house” (i.e. the game organizer). The paper then generalizes this intuition for real-life situations and suggests ways to improve our ability to profit from risky investment opportunities.[4]

Keywords: Decision Making under Uncertainty, Risk, Uncertainty, Utility, Risk Aversion, Kelly, Gambling, Betting, Coin Flip, Market Anomalies, Return Predictability, Behavioral Finance, Market Timing, Gamblers Fallacy, St Petersburg Paradox, CRRA, Mathematical Methods, Operations Research, Statistical Decisions Theory, Asset Pricing, Programming Models

JEL Classification: B12, B16, B20, C00, C02, C10, C11, C44, C50, C57, C61, C63, C73, D03, D81, D83, E00, G00, G02, G11, G12, G14, G17, G23

---

[1] The author is grateful to Victor Haghani, Richard Dewey, Arjun Viswanathan and Aaron Brown for helpful comments, suggestions and feedback during the work on this paper. All mistakes are, of course, my own.

[2] Since starting with LTCM in 1994 as a quantitative strategist Vladimir Ragulin (Vladimir.Ragulin@gmail.com), has worked in trading positions with both hedge funds and banks, focusing on quantitative trading strategies and risk management.

[3] Haghani, Victor and Dewey, Richard, Rational Decision-Making under Uncertainty: Observed Betting Patterns on a Biased Coin (October 19, 2016). Available at SSRN: https://ssrn.com/abstract=2856963

Introduction

This research note is inspired by a simple but illuminating behavioural finance experiment by Victor Haghani and Richard Dewey of Elm Partners[5], built around flipping for profit an unfair coin with known heads/tails probabilities (60/40). Using this elementary framework with fully known game parameters and well-understood win/loss probabilities the paper builds a simple model showing that our ability to take advantage of this clearly profitable opportunity greatly depends on our attitude to risk, our outside financial resources as well as on financial strength of "the house" (i.e. the game organizer). The author hopes that this paper will make readers better risk-takers and investors, highlighting the importance of capital, courage and discipline in achieving financial goals, and showing that some games are much better than others in their ability to generate wealth.

Haghani and Dewey offered a group of financial professionals and economics students a simple money-making opportunity. Starting with a $25 stake the players could make, over the next 30 minutes, repeated bets on an unfair coin biased to come up heads 60% of the time. The participants could make up to 300 flips[6]. On each bet a player could bet on either heads or tails up to the full amount of her capital at the time. Those players that did not go bust by the end of the 300 flips would get paid their final winnings, subject to a cap of $250 (which was revealed to a participant once she had a chance of exceeding the cap upon winning the next bet).

Despite an extremely simple game set up and considerable academic training by the participants the results were shockingly disappointing as "suboptimal betting come in all shapes and sizes: over-betting, under-betting, erratic betting and betting on tails"[7]. Even though the game is heavily stacked in favour of the participants with the expected payoff of just over $240 and a 95% chance of being able to reach the $250 maximum, about 28% of the subjects actually went bust, losing their full $25 stake. Even those subjects who did not go bust performed well below potential with only 21% reaching the maximum and the remaining 51% averaging an average pay out of $75.

After having delivered a decidedly sub-par performers during the "live play", the author was motivated, after the fact, to solve for the optimal game strategy. A numerical approach, using a lattice solver, is described below. As this paper was being written, two researchers, Aaron Brown[8] and Arjun Viswanathan[9], published analytical solutions for this game's optimal strategy, which are consistent with my results. Both Brown and Viswanathan use recombining trees and apply their frameworks to more general classes of games. Brown explores games with an unknown cap as well as proposes a number of simple but effective heuristics, which allow players to capture a significant proportion of the value while avoiding the need to find an exact solution. Viswanathan considers players with outside wealth and derives a tree-based solution for a general class of utility function.

The contribution of this paper is twofold. First, unlike Brown and Viswanathan, I apply methods of classical financial economics (portfolio theory and Bellman equation) to solve the problem numerically. The use of well-known economic tools makes it easier to see connections between the game and various real-life investment problems and to develop economic intuition. Second, rather than focusing on the best way to play the game, the paper focuses on the game value in alternative "real-life" scenarios, i.e. how much different individuals with alternative preferences and outside wealth would be willing to pay to play the game. The paper establishes several insights

[5] Haghani and Dewey.
[6] However, the authors estimated that the average number of flips was around 100.
[7] Haghani and Dewey, p.3.
[8] Brown, Aaron, Optimal Betting Strategy with Uncertain Pay-out and Opportunity Limits (December 8, 2016).
[9] Viswanathan, Arjun, Generalizing the Kelly strategy (December 6, 2016). Available at arXiv:1611.09130v3.

into the relationship between individual risk preferences, capital and their ability to take advantage of profitable but risky investment opportunities.

Base Case Solution

We start by deriving the optimal strategy for the original Haghani-Dewey flipping game. The game can be thought as a simplified version of a standard multi-period portfolio optimization problem as described, for example, in Merton[10] or Campbell[11] (with our coin flip thought as an investment in a risky asset), which is solved using the Bellman Equation technique[12]. We define a "game value" function $J(S,n)$ that is equal to the maximum expected gain, assuming optimal play, for a player that has a stake of S with n coin flips left to go. Those familiar with multi-period portfolio optimization literature[13] will no doubt notice that our "game value" function is exactly equivalent to the derived utility or wealth. At the end of the game (i.e. at the boundary) it's clear that $J(S,0) = \min(S,C)$, where C is the payout cap. Then for n>0 we can solve for $J(W,n)$ backwards, using a recursive formula:

$$J(S,n) = \max_{0\leq b\leq W}[0.6 * J(S+b,n-1) + 0.4 * J(S-b,n-1)] \qquad (1)$$

where we optimize over the size of the bet b. Since the boundary condition is concave and the optimization interval is bounded, the maximum will always exist, but, as Brown and Viswanathan show in their separate analytical derivations, it will generally not happen for a unique optimal bet b.

The boundary condition and the recursive relationship (1) is sufficient to solve for the game value function numerically for any values of player stake S or number of flips remaining n using an explicit lattice solver.

The solver also produces, for each point of the lattice, an optimal bet value b, which, unfortunately need not be unique (as shown by Brown and Viswanathan), but since this paper's focus is the game value, rather than the optimal strategy, it is not a serious problem for my purposes. For uniformity, the solver was calibrated to take the highest bet level that maximized the game value at each lattice point.

Chart 1 summarizes the results, showing the game value $J(S,n)$ on the vertical axis against the player's stake S on the horizontal axis, and each line corresponding to a different number of flips remaining ($n = 1, 50, 100, 300$).

---

[10] Merton, Robert C, Optimum Consumption and Portfolio Rules in a Continuous-Time Model (September 30, 1970), Journal of Economic Theory 3, 373-413.
[11] Campbell, John Y. and Viciera, Luis M., Strategic Asset Allocation: Portfolio Choices for Long-Term Investors. Oxford University Press, 2002.
[12] Bellman, R.E. 1957. *Dynamic Programming*. Princeton University Press, Princeton, NJ. Republished 2003: Dover, ISBN 0-486-42809-5.
[13] Merton or Campbell and Viciera in Bibliography.

Chart 1: Expected Value of the Game vs. Player's Initial Stake

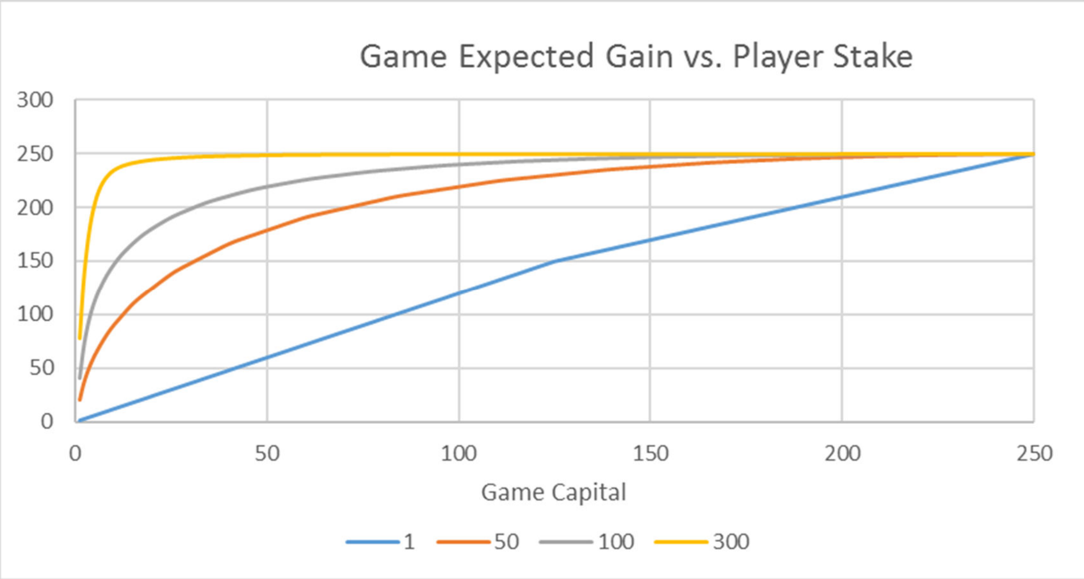

It is reassuring that for n=300 my game value of $246 matches the analytical solutions by Brown and Viswanathan.

Rules of Thumb -- Good and Bad

Having solved the problem by "brute force", I was curious if there were some simple rules of thumb or heuristics which would allow a player to get close to the optimal result using a simple mental rule. From discussions with Haghani and Dewey it appeared that two simple heuristics were common among their subjects. Both heuristics ignored the pay-out cap, which is somewhat understandable since the cap was only revealed later in the game only if a participant had a chance of exceeding it.

The first heuristic ("uncapped expected gain") correctly observed that in the absence of a cap each bet's expected gain monotonically increases with the bet size, and, hence, bet the entire capital on every flip. If we make an additional assumption that a player is risk-neutral and only cares about her the expected payoff and not the expected utility, then such strategy would indeed be optimal. However, these assumptions do not apply in our experiment and this turns out to be a very bad rule of thumb, indeed, offering only a 13% probability of winning the $250 maximum and an 87% of going bust, for an expected gain of just over $32.

The second heuristic ("ratio betting") bets a fixed ratio of capital on every flip until the cap is revealed, and then reduces the bet so as not to exceed the maximum pay-out. Haghani and Dewey highlight the "Kelly" rule, which recommends betting on every flip the ratio equal to the percentage edge of the player, i.e. 60% - 40% = 20% of capital. It is a well-known result that betting this ratio maximises the expected return on the player's capital. Even though the Kelly rule also does not consider the pay-out cap until much later in the game, it happens to be a very good heuristic, offering an expected gain of around $238, according to Brown, i.e. within 4% of the absolute theoretical maximum.

The two heuristics above can be viewed as myopic rules that only look 1 flip ahead and, for every flip, maximize the expected utility of an investor with a specific risk-aversion schedule – the first rule for a risk-neutral investor ("linear utility") and the second for a risk-averse investor with log utility. [14]

These observations motivated the author to employ the coin-flipping framework to gain more general insights about a broader class of investment games with positive expected returns and the ability of players with different risk preferences to make money playing such games.

Playing Against Deep Pockets – Changing the Pay-out Cap

A closer look at the Chart 1 suggests that the game value depends just as much on the cap (i.e. maximum pay-out) as on the player's starting capital. In fact, for a game with n=300 flips, the expected value gets to within 10% of the maximum cap for any starting stake greater than $8. After noting that this is yet another illustration of how attractive the 60/40 proposition was to the players (with the main binding constraint on their winnings being the $250 maximum cap), I used the numerical solver to explore how the game value and betting strategy changes as we increase the cap. Results (shown in Charts 2a – 2c) confirm the tremendous power of compounding, even in the presence of uncertainty. With a $1mm cap the value of the game starting with a $25 stake is more than $277,000, and if the cap is $10mm, the expected gain is over $1.3mm. The expected value of an uncapped game is an astronomical $\$25 * 1.2^{300} = 1.42 * 10^{25}$. Indeed, it was wise for Haghani and Dewey to impose a winnings cap!

Chart 2a

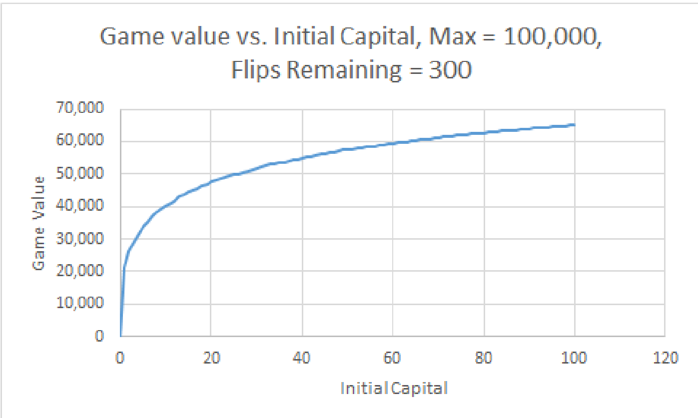

[14] For a more detailed exploration see Haghani, Victor and Morton, Andrew, Optimal Trade Sizing in a Game with Favourable Odds: The Stock Market (November 25, 2016). Available at SSRN: https://ssrn.com/abstract=2875682.

Chart 2b

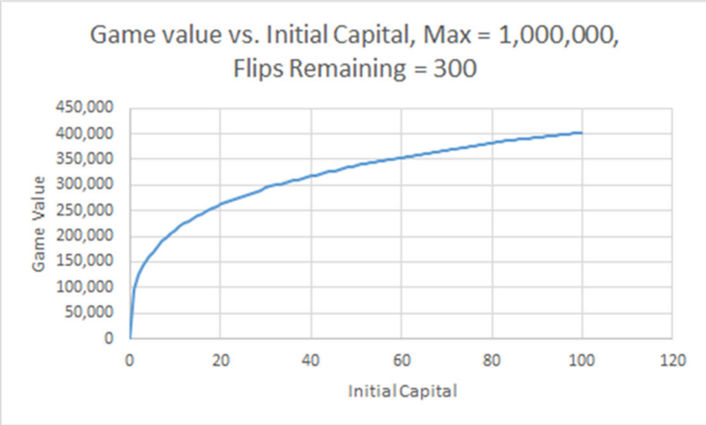

Chart 2c

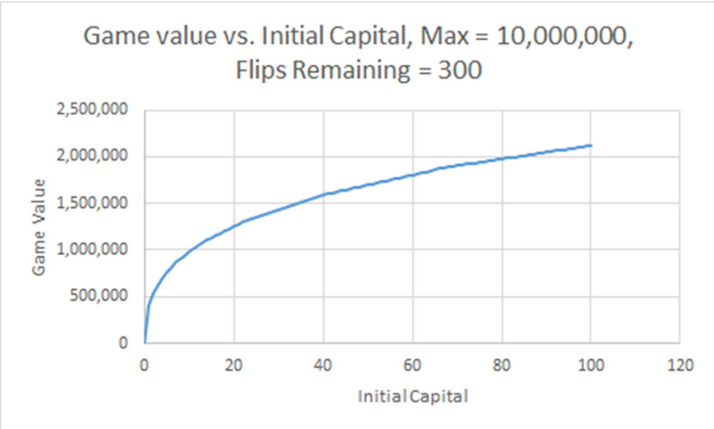

Alternative Player Risk Preferences

Audentes fortuna iuvat ("Fortune favors the bold", Virgil, Aeneid)

Once we know the optimal strategy, are there any obstacles to monetizing the potential gains offered by the coin-flipping opportunity, especially for larger values of the cap? It is easy to see that the distribution of final winnings, especially for greater cap values is heavily skewed, with significant fraction of expected value coming from very unlikely extremely high pay-outs. Since Bernoulli's St. Petersburg Paradox, economists have long known[15] that risk-averse individuals would value such uncertain payoffs at a discount, often significant, to their mathematical expectation. In this section, we will explicitly consider alternative player risk preferences and show how the game strategy needs to be modified to take risk-aversion into account.

It is easy to incorporate players' risk preference by slightly modifying our lattice solver to maximize expected utility of the game rather than expected monetary payoff. To make the model even more realistic we can allow outside wealth to describe behaviour of players to whom the loss of the initial $25 stake is inconsequential, or generally address a common concern that many standard utility functions assign infinitely negative utility to zero wealth, implying excessively cautious betting at low levels of game capital,

Specifically, assume that a final pay-out of S at the end of the game gives the player a utility of $U(S;W)$, where W represents the player's wealth outside of the game. The outside wealth W is constant during the game (so that the player's total wealth is $(S+W)$.

We can then re-define the game value $J(S,n;W)$ as the maximum expected *utility* a player can possibly achieve with a stake of S and n flips remaining. The boundary condition becomes $J(S,0;W) = U(\min(S,C)\,;W)$, while the backward induction formula remains the same with the only addition of dependency on W, the investor's outside wealth.

$$J(S,n;W) = \max_{0\le b\le W}[0.6 * J(S+b,n-1;W) + 0.4 * J(S-b,n-1;W)] \quad (2)$$

We consider 3 alternative player risk preferences:

1. Linear Utility (i.e. Risk-Neutral): $U(S;W) = S + \mathrm{W}$.
2. Log Utility: $U(S;W) = \ln(S+W)$.
3. Constant Relative Risk Aversion Utility (CRRA(2)): $U(S;W) = \frac{(S+W)^{1-\eta}-1}{1-\eta}$, with $\eta = 2$. This implies an investor twice as risk-averse as a log utility investor.

Clearly, player 1 is risk-neutral (i.e. only cares about the game's expected payoff) and we have already solved this case in the previous sections. Players 2 and 3 are risk-averse and would value uncertain payoffs at less than their mathematical expected values. To solve the game for risk-averse players, we need to know how much wealth they have outside of the game, i.e. how painful for their lifestyle it would be to lose all the game capital. On the other hand, risk-neutral player 1 would play the same way regardless of her outside wealth.

[15] The difference between mathematical expectation of a risky payoff and the value that an individual would pay for such uncertain bet has been illustrated, for example, by Bernoulli's St. Petersburg Paradox in 1738.

Case 1: No Outside Wealth

For a risk-averse player with no outside wealth (W = 0) it would be extremely painful to lose all the money in the game (and end up with a negatively infinite utility), therefore, we would expect her to only bet a fraction of capital to minimize chances of going bankrupt. We know from basic utility theory that in the absence of a cap the Log-Utility player 2 would bet in line with the Kelly formula (i.e. 20% of capital), the CRRA(2) player would bet half of Kelly (i.e. 10%).[16]

Table 1 and Chart 3 show how much each player would pay to play the coin-flipping game with different caps, i.e. her Certainty Equivalent Value. It is not a surprise that more risk-averse players would pay less, in some cases by a lot, for this clearly profitable opportunity. In fact, the CRRA(2) player would bet so conservatively that his game value does not materially increase (unlike other players) as the maximum pay-out cap is raised from $100k to $10mm.[17]

Table 1: Game Certainty Equivalent Value vs. Max Payout

Max Payout\Utility	Linear	Log	CRRA(2)
100,000	50,148	7,271	526
1,000,000	277,686	9,562	526
10,000,000	1,338,617	10,264	526

Chart 3

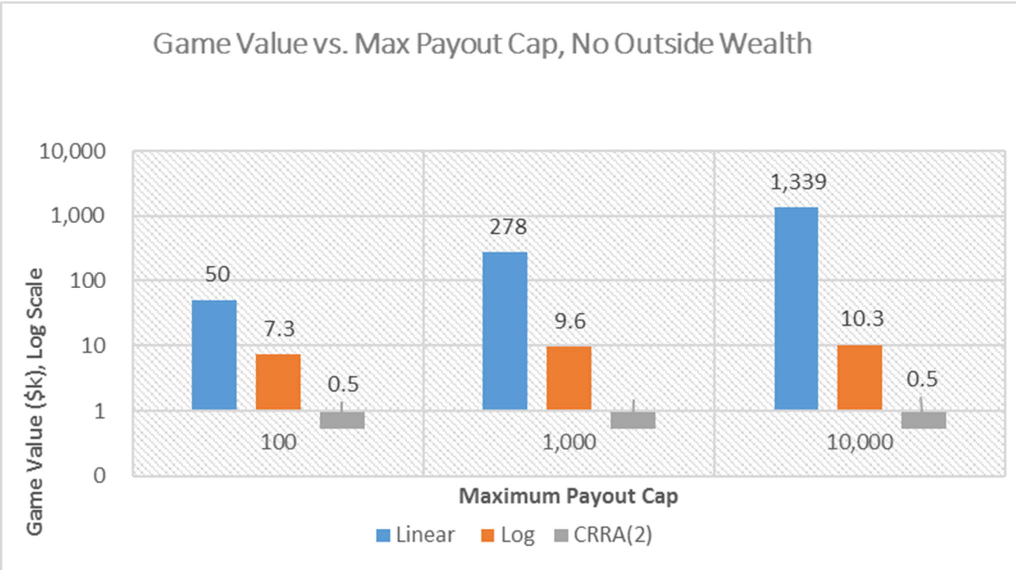

Digging a little deeper, there are 2 reasons why risk-averse players would pay less to participate in the game:

1. Since they are more concerned about losing money, they tend to bet less aggressively and, therefore, do not fully "monetize" the positive edge offered by the opportunity.
2. Even if bets and payoffs were the same, the more risk-averse players would have a lower Certainty Equivalent Value for the same distribution of outcomes.

[16] Campbell and Viciera.

[17] The Game Value for the CRRA(2) player does grow, but by less than the rounding error, which can also be seen from increasing Expected Payoff in Table 2. But the CRRA(2) player does not value the unlikely probability of very big wins as highly as other players.

Table 2 illustrates relative importance of these 2 reasons, showing the Expected Payoff (i.e. before we apply Utility Functions) for each of the players, when they bet in a way that's optimal for them. In other words, the Expected Payoff is what a risk-neutral Player 1 would pay to play the game with the same strategy as the risk-averse player. We can see that differences in Expected Payoffs are still significant, but a lot smaller than differences in valuations. Hence, risk-averse players 2 and 3 both make less money due to more conservative betting and, furthermore, apply greater discount to those risky cash flows, resulting in dramatic differences in valuations of the opportunity.

Table 2: Game Expected Payout vs. Max			
Max Payout\Utility	Linear	Log	CRRA(2)
100,000	50,148	38,499	8,403
1,000,000	277,686	162,216	10,087
10,000,000	1,338,617	485,382	10,090

Case 2: Players with Outside Wealth

In discussions of the Haghani-Dewey experiment a common defence for so many participants who went bust was that since the $25 initial stake did not represent a meaningful portion of their wealth (after all, many were finance professionals!), they could afford to act as if they are risk-neutral and just bet the maximum amount where they had a favourable edge.

Ignoring for the moment that betting the maximum is only optimal when there is no cap, we will now consider how outside wealth changes players' valuation of the game and their betting strategies.

It's easy to see that for a risk-neutral player outside wealth makes no difference, while for a risk averse player having outside resources helps in 2 ways already discussed above. First, a wealthy risk-averse player is less concerned about losing 100% of his game stake, therefore he will bet more aggressively and generally will be able to better "harvest" the positive edge offered by the unfair coin (assuming he plays optimally). Second, a wealthy risk-averse player would pay more (in terms of her Certainty Equivalent) for a given risky payoff distribution than a player with the same utility function but without wealth.

Below in Tables 3 and 4 as well as Chart 4 we show game valuations and expected payoffs for players with different utility functions and $1mm wealth outside of the game. (I.e. outside wealth is the only difference between Tables 1 and 3, and Tables 2 and 4).

Chart 4

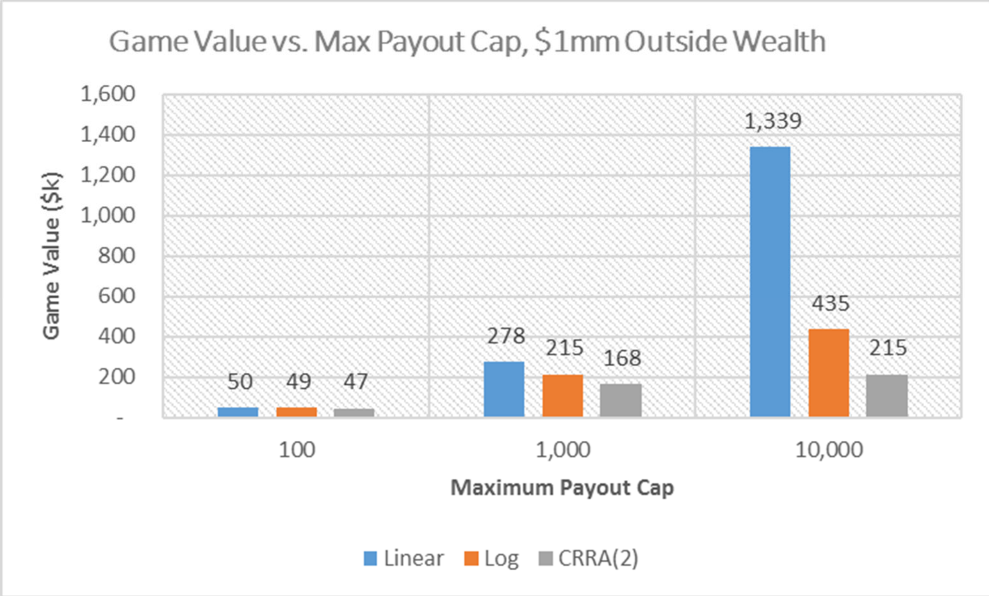

Table 3: Game Certainty Equivalent Value vs. Max Payout			
Max Payout\Utility	Linear	Log	CRRA(2)
100,000	50,148	48,548	47,413
1,000,000	277,686	215,287	167,611
10,000,000	1,338,617	434,931	214,651

Table 4: Game Expected Payout vs. Max Payout			
Max Payout\Utility	Linear	Log	CRRA(2)
100,000	50,148	49,685	49,685
1,000,000	277,686	277,885	271,345
10,000,000	1,338,617	1,222,850	901,233

As expected, players with outside wealth both achieve higher expected payoffs and value these payoffs higher. The difference is especially stark for the CRRA(2) player, for whom the valuation of the game with a $10mm cap grows from $526 to $214,651! However, even for wealthy players, their ability to take advantage of the opportunity varies greatly depending on their attitude to risk, especially for game with higher caps. For example, the risk-neutral player would pay almost 6.25 times as much to play the game with a $10mm cap than the CRRA(2) player.

Conclusion and Topics for Future Research

The paper develops a numerical lattice-based approach to solving the Haghani-Dewey coin-flipping game and generalizes it further, looking at various pay-out caps and considering players with different attitudes to and financial resources. It is reassuring that the numerical solution gives similar results to analytic solutions by Brown [2] (in the risk-neutral case) and Viswanathan [3] (for a general class of utility functions). The value of this paper's numerical approach, however, is in its flexibility, which allows easy changes in parameter and game design to generate comparative statics and provide important economic intuition for more real-life settings. Furthermore, unlike previous works, this paper identifies two separate effects that determine a player's ability to monetize this profitable opportunity, looking separately at the changes in the expected payoff distribution and on how the player's attitude to risk impacts her valuation of this uncertain payoff.

One of the motivations for the Haghani-Dewey experiment was to provide intuition about risk-taking and optimal investment strategies in real-life settings. They showed that a simple Kelly heuristic dramatically outperformed most of the participants and argue that, while not strictly optimal for this experiment, a Kelly-type approach would be beneficial to many investors in their financial decision-making. This paper provides further practical examples of how this intuition can be applied by individuals in making decisions under uncertainty. For example, it is well known that the Kelly Rule is mathematically optimal for investors with Log Utility playing games with no upper limit. At the same time, studies in behavioural finance [6] show that most investors display more risk aversion than implied by Log Utility, hence a common recipe is to adjust the Kelly rule in various ways, for example, by using a smaller betting ratio. This paper quantifies economic consequences of such adjustments, which can be very significant – up to several orders of magnitude in terms of the player's expected gain.

In other words, dramatically different valuations of the simple Haghani-Dewey game for players with different utility functions show the importance of an individual's attitude to risk in her ability to take advantage of economic opportunities. Despite its importance, individual risk-aversion is rarely chosen based on a rational economic calculation. Often it is influenced by the culture, institutional settings, or recent experiences (e.g. recent financial crisis).

This paper shows, using a comparison of bettors with different levels of financial wealth, that we should be much more deliberate and thoughtful in choosing our attitudes to risk, carefully considering our resources, commitments, aspirations and trade-off preferences between security and chances of success. Without such analysis, a simple use of "socially representative" risk tolerance, can either lead to an individual falling far short of maximum potential or financially over-extended.

The author hopes that this paper contributes to our understanding of the relationship between individual risk appetite and economic outcomes and help individuals be more rational and disciplined in choosing their personal "risk limits" to make best use of potentially profitable but uncertain projects they may encounter in their own business activities.

Bibliography:

1. Haghani, Victor and Dewey, Richard, Rational Decision-Making under Uncertainty: Observed Betting Patterns on a Biased Coin (October 19, 2016). Available at SSRN: https://ssrn.com/abstract=2856963 .

2. Brown, Aaron, Optimal Betting Strategy with Uncertain Pay-out and Opportunity Limits (December 8, 2016). Available at SSRN: https://ssrn.com/abstract=2947863 or http://dx.doi.org/10.2139/ssrn.2947863 .

3. Viswanathan, Arjun, Generalizing the Kelly strategy (December 6, 2016). Available at arXiv:1611.09130v3 .

4. Haghani, Victor and Morton, Andrew, Optimal Trade Sizing in a Game with Favourable Odds: The Stock Market (November 25, 2016). Available at SSRN: https://ssrn.com/abstract=2875682.

5. Merton, Robert C, Optimum Consumption and Portfolio Rules in a Continuous-Time Model (September 30, 1970), Journal of Economic Theory 3, 373-413.

6. Campbell, John Y. and Viciera, Luis M., Strategic Asset Allocation: Portfolio Choices for Long-Term Investors. Oxford University Press, 2002.

7. Bellman, R.E. 1957. *Dynamic Programming*. Princeton University Press, Princeton, NJ. Republished 2003: Dover, ISBN 0-486-42809-5.