Dynamic Gambling under Loss Aversion*

Yair Antler

University of Essex

November 22, 2017

Abstract

A loss-averse gambler faces an infinite sequence of identical unfair lotteries and decides in which of these lotteries to participate, if at all. We establish that it is always possible to find an unfair baseline lottery such that the gambler chooses to participate in several such lotteries, and that the gambler's optimal gambling plan is a left-skewed stopping rule. We introduce stochastic updates to the gambler's reference point and show that these updates create dynamic inconsistencies in the gambler's behavior. We find that dynamically inconsistent sophisticated gamblers participate in fewer lotteries than dynamically inconsistent naive ones, and that dynamically inconsistent naive gamblers may participate in fewer lotteries than dynamically consistent ones.

*This paper is based on one of the chapters of my PhD dissertation. I am highly indebted to Rani Spiegler for his encouragement and many insightful discussions. I acknowledge financial support from ERC grant no. 230251. I thank Ayala Arad, Eddie Dekel, Kfir Eliaz, and Asher Wolinsky for helpful conversations and comments. All remaining errors are mine. *Correspondence:* Yair Antler, University of Essex. *E-mail:* yair.an@gmail.com.

1 Introduction

Casino gambling has social and economic effects both on communities and on gamblers at the individual level. In the US, in 2016, the tribal and commercial casino industries generated a combined total annual gross gaming revenue of more than 70 billion dollars (American Gaming Association, 2017) and over 81 million individuals visited a casino. The existence of individuals who pay insurance premiums to reduce their exposure to risk and at the same time participate in gambling activities has intrigued economists since Friedman and Savage's (1948) seminal work. Prominent explanations for this phenomenon include utility from gambling (Conlisk, 1993) and probability distortion à la commutative prospect theory (Barberis, 2012).

Loss aversion is *a tendency to evaluate changes or differences rather than absolute magnitudes* and *to dislike losses more than comparably sized profits.* Loss aversion is one of the most well-documented phenomena in economics and psychology. The foundations for reference-dependent preferences were laid by Markowitz (1952), and loss aversion as we know it nowadays was introduced by Kahneman and Tversky (1979). Loss-averse economic agents are typically averse to most risks. For example, a loss-averse agent always rejects a single 50:50 fair lottery since the pain he would suffer from the potential loss is greater than the pleasure he would derive from the potential profit.

We examine the behavior of a loss-averse gambler who faces an infinite sequence of unfair lotteries, each of which pays 1 with probability $p < 0.5$ and -1 with probability $1 - p$. The gambler has to decide whether or not to participate in each of these lotteries (he can participate in as many lotteries as he wishes). Observe that since each lottery is unfair and the gambler dislikes losses more than he likes gains of the same magnitude, participation in a fixed number of $k > 0$ lotteries is unattractive to the gambler (i.e., the gambler prefers not participating at all to participating in k lotteries). Will the gambler refuse to play in the above lotteries or will he choose to participate and implement a "more complex" gambling plan? Can a casino that offers such lotteries affect the gambler's decision?

Even though the gambler is loss averse and the lotteries are unfair, he may decide to participate in several lotteries. In such cases, his optimal gambling plan is a left-skewed stopping rule (i.e., the gambler stops after accumulating profits of $h > 0$ or losses of $-l$, where $l > h$). We show that for every set of preference parameters, it is possible to find $p < 0.5$ such that the gambler will participate in these lotteries. This implies that a casino that offers such lotteries and controls p can make a strictly positive profit when it faces loss-averse gamblers.

In practice, casinos often offer monetary benefits or other perks to incentivize gamblers to bet. Typically, they incur short-term losses on these offers but gain the losses back as the gamblers continue gambling. Can a casino that offers the above lotteries benefit from offering monetary benefits to individual gamblers? In particular, suppose that the casino can offer a transfer $\tau > 0$, and accepting this transfer obliges the gambler to participate in one lottery (after the lottery is realized, the gambler is allowed to stop or continue gambling as much as he wishes). Can such a transfer make both the gambler and the casino better off? Observe that such agreements between a risk-neutral casino and a *risk-averse gambler* are not viable.[1] However, we show that for a wide range of values of p, such an agreement makes both the casino and a loss-averse gambler better off.

The gambler's preferences are defined over gains and losses with respect to a reference point (e.g., the gambler's wealth at the beginning of the game). Changes in the reference point may affect the gambler's preferences and may result in *dynamically inconsistent* behavior as the gambler's decision at a specific wealth level depends on his reference wealth. We extend our model by allowing for *stochastic updates to the gambler's reference point*: we assume that in every period, with some probability, the gambler's reference point is updated to the gambler's current wealth level. We interpret such an update as internalizing the gambler's profits (or losses) since the last update. The gambler may or may not be aware of the possibility of future updates. We shall refer to a gambler who is aware of these updates as *sophisticated* and to a gambler who is unaware of these updates as *naive*. To analyze the effect of these updates we apply a multi-selves approach (Strotz, 1956). We establish that a sophisticated gambler, in expectation, participates in fewer lotteries than a naive one. Surprisingly, a dynamically inconsistent naive gambler may participate in fewer lotteries than a dynamically consistent gambler (i.e., a gambler whose reference point is never updated).

The paper proceeds as follows. We present the model in Section 2 and analyze it in Section 3. In Section 4 we extend the model by allowing for updates to the gambler's reference point. Section 5 covers related literature and Section 6 concludes. All proofs are to be found in the Appendix.

[1] A risk-averse gambler accepts a risky prospect only if its expected value is strictly positive (i.e., if, including the transfer τ, the casino's expected profit is strictly negative).

2 The Model

A gambler faces an infinite sequence of identical unfair lotteries, each of which pays 1 with probability $p < 0.5$ and -1 with probability $1-p$. There is a discount factor $\delta < 1$. At each time $t = 1, 2, 3, ...$, the gambler decides whether or not to participate in a lottery. Let a_t denote the gambler's decision at time t and let w_t denote his wealth at the beginning of time t. Let r_t denote the gambler's *reference wealth* at time t. We assume that $r_1 = w_1$ (i.e., the gambler's initial reference point is his wealth before he starts gambling). Let $x_t := w_t - r_t$ denote the gambler's *gains* or *losses* with respect to his reference wealth at the beginning of time t.

The gambler's preferences are defined over gains and losses with respect to his reference wealth. They are represented by

$$U(x) = \left\{ \begin{array}{ccc} u(x) & \text{if} & x \geq 0 \\ -v(-x) & \text{if} & x < 0 \end{array} \right\}$$

where $u : \mathbb{R} \rightarrow \mathbb{R}$ and $v : \mathbb{R} \rightarrow \mathbb{R}$ are strictly increasing, unbounded, and satisfy $u(0) = v(0) = 0$. We assume that $u(x) < v(x)$ for every $x > 0$, that $\frac{v(x)}{x}$ is weakly decreasing in x, and that $lim_{x \rightarrow \infty} \frac{v(x)}{x} = 0$. In words, we assume that the gambler is loss averse, and that his sensitivity to gains or losses diminishes as these gains or losses increase.[2] These assumptions are satisfied by the following utility function, proposed by Kahneman and Tversky (1992):

$$U(x) = \left\{ \begin{array}{ccc} x^{\alpha} & \text{if} & x \geq 0 \\ -\lambda(-x)^{\beta} & \text{if} & x < 0 \end{array} \right\} \tag{1}$$

where $\lambda > 1$, and $0 < \alpha \leq \beta < 1$.

Denote the history at time t, $(r_1, x_1, a_1, ..., r_{t-1}, x_{t-1}, a_{t-1}, r_t, x_t)$, by h_t. A strategy a maps the history at time t to a decision whether or not to participate in a lottery at time t. We shall denote by $V(a, x, p, \delta)$ the gambler's expected value from following the strategy a when his current gains level is x. A gambling strategy is said to be *stationary* if, for every time t, the decision a_t is conditioned only on x_t. By Theorem 7 in Blackwell (1965), there exists a stationary strategy that maximizes the gambler's expected payoff. In the present setting, there exists a *non-randomized stationary strategy* that maximizes the gambler's payoff. To see this, observe that if the gambler uses a randomized stationary strategy α, then

[2] We allow for some segments in which u and v are not concave.

at each gains level x at which he is scheduled to randomize, he obtains a value of $V(x, a, p, \delta) = U(x)$. Conditional on reaching gains level x, the gambler can obtain $U(x)$ by stopping immediately. Thus, by switching the gambler's decision in every node in which he is scheduled to randomize, we can obtain a pure strategy that guarantees the gambler the same payoff as the one he attains under[3] α. We shall restrict our attention to non-randomized stationary strategies.

3 Optimal Gambling

In this section, we characterize the gambler's behavior. Lemma 1 establishes that a strategy is optimal only if it is a stopping rule. That is, the gambler plays until either he accumulates gains of $h \geq 0$ or losses of $-l \leq 0$. When $h = l = 0$, it means that the gambler does not gamble. We shall refer to such a strategy/stopping rule as a *degenerate* one.

Lemma 1 *An optimal strategy must induce a stopping rule: the gambler stops participating in lotteries whenever his gains level reaches $h \geq 0$ or $-l \leq 0$.*

We shall now use the technical result of Lemma 1 in order to examine the gambler's optimal gambling strategy further. We refer to a stopping rule under which the gambler stops after accumulating gains of $h > 0$ or after accumulating losses of $-l < 0$ as a *left-skewed* one if $l > h$. The next result establishes that the gambler's optimal stopping rule[4] is either degenerate or left-skewed.

Proposition 1 *The optimal gambling strategy is either degenerate or a left-skewed stopping rule.*

Left-skewed stopping rules are attractive to the gambler since they provide him with more opportunities to *break even* after accumulating some losses. He values these opportunities since he is risk-seeking in (some parts of) the losses segment of his utility function.

A natural question that arises is whether the gambler starts gambling at all. The next result establishes that for large values of $p < 0.5$, the optimal gambling plan is not degenerate.

Proposition 2 *There exist $p^\star < 0.5$ and $\delta^\star < 1$ such that if both $p \in (p^\star, 0.5)$ and $\delta \in (\delta^\star, 1)$, then the gambler's optimal strategy is a non-degenerate left-skewed stopping rule (i.e., the gambler participates in several lotteries).*

[3] This is not a general observation. For example, Henderson et al. (2017) show that a gambler who distorts probabilities may obtain a strictly higher payoff by means of a randomized strategy.

[4] Generically, the gambler's optimal strategy is unique as small changes in p break his indifference between different stopping rules. When the gambler is indifferent between several stopping rules, each of them is either degenerate or left-skewed.

Proposition 2 establishes that facing a sufficiently large $p < 0.5$, the gambler will gamble. Note that since the gambler is loss averse, he does not want to participate in any *fixed number* of k lotteries (as even for $p = 0.5$ such a strategy induces a combination of 50:50 fair lotteries, which are unattractive from a loss-averse gambler's perspective; see (6) in the proof of Lemma 1). However, this does not imply that the gambler is risk averse. His diminishing sensitivity to losses makes him risk-seeking in some parts of the losses segment of his utility function. Left-skewed stopping rules allow the gambler to enjoy more gambling at gains levels at which he is risk-seeking while not gambling so much at gains levels at which he is risk averse.

A natural question to ask is how large must p be for the gambler to participate. This, of course, depends on the gambler's preferences. For example, consider the preferences represented by (1). The larger λ and β are, the larger the cutoff p is. Set $\lambda = 1.1, \alpha = \beta = 0.87$, and consider the $\delta = 1$ limit. Observe that since the gambler is dynamically consistent, his optimal stopping rule must maximize his expected value when $x = 0$. The problem is an immediate application of the "gambler's ruin" problem (for a textbook treatment, see, Grinstead and Snell, 1997). Thus, we need to choose h and l that maximize:

$$lim_{\delta \to 1} V\left(a^{\star}, 0, p, \delta\right) = \frac{1 - \left(\frac{1-p}{p}\right)^{l}}{1 - \left(\frac{1-p}{p}\right)^{l+h}} h^{\alpha} - \lambda \frac{1 - \left(\frac{p}{1-p}\right)^{h}}{1 - \left(\frac{p}{1-p}\right)^{l+h}} h^{\beta} \tag{2}$$

If $p = 0.49$, then the optimal gambling plan is to stop after accumulating gains of $h = 1$ or after accumulating losses of $-l = -7$. This strategy induces a winning probability of 0.857 (i.e., the probability of finishing the game with $x > 0$ is 0.857) and a strictly positive expected value for the gambler. Given this strategy and these parameters, a casino that offers such lotteries makes an expected profit of 0.146 (i.e., the gambler's expected loss is 0.146).

Example: Gambling Inducements

Our results in the previous section established that a casino that offers the above lotteries and can control the baseline lottery's probability p, is able to make a positive expected profit at the expense of loss-averse gamblers. In practice, however, it is not always possible to fully adjust p as gambling games are often canonic games (e.g., Blackjack) with given probabilities and gamblers may be unwilling to play new games whose rules they are unfamiliar with (e.g., they might not understand the rules of these games). What can a casino do when it cannot control p?

We shall now demonstrate that a casino that offers the above lotteries can benefit from providing compensation to gamblers who start gambling. Consider the following interaction between a risk-neutral casino and our loss-averse gambler. The casino commits to a potential monetary transfer τ to the gambler. The transfer is conditioned on gambling in at least one lottery (say, at time $t = 1$). If the gambler participates in that lottery, then he receives τ (after participating in one lottery, he is allowed to participate in as many lotteries as he wishes). We interpret the compensation τ as benefits available to gamblers who enter the casino.

Initially, it is unclear whether or not a casino can benefit from such an offer. For example, a casino would never pay to incentivize a risk-averse gambler to enter the casino. This is because a risk-averse gambler would stop gambling after the first lottery and a compensation greater than $1-2p$ (i.e., the casino's expected profit) would be required to make him start gambling, as risk-averse economic agents find prospects with negative expected values unattractive. We will show that if p is not too large[5] or too small, the casino can benefit from offering a transfer $\tau > 0$ to gamblers who start betting (i.e., in return for participating in at least one lottery).

Consider the $\delta = 1$ limit and the preferences represented by (1). Suppose that $p = \frac{18}{37}$ as in the Red or Black roulette game, and let $\alpha = \beta = 0.8$ and $\lambda = 1.2$. If the gambler stops gambling after accumulating gains of h or losses of l, then $lim_{\delta \to 1} V(a, 0, \delta, p) < 0$ as there exist no two integers h and l such that

$$\frac{1 - \left(\frac{1-p}{p}\right)^l}{1 - \left(\frac{1-p}{p}\right)^{h+l}} h^{0.8} - 1.2 \frac{1 - \left(\frac{p}{1-p}\right)^h}{1 - \left(\frac{p}{1-p}\right)^{h+l}} l^{0.8} \geq 0 \tag{3}$$

It follows that the gambler's optimal stopping rule is degenerate (i.e., $h = l = 0$). Hence, the casino's expected profit is 0 if it does not incentivize the gambler to start betting. We now illustrate how both parties can benefit when the casino makes a transfer $\tau^\star = 0.01$ to the gambler, who, in return, starts gambling.

Given a transfer of $\tau^\star = 0.01$, the following strategy $a^\star$ induces a strictly positive expected value for the gambler at his reference wealth[6]: stop once you

---

[5]By Proposition 2, for values of p sufficiently close to 0.5, the gambler will participate regardless of whether or not he is offered compensation in return. A casino has no reason to compensate the gambler in such a case.

[6]For completeness, the gambler's expected payoff $V(a^\star, 0, p, \delta)$ is at least $\frac{1-\left(\frac{1-p}{p}\right)^8}{1-\left(\frac{1-p}{p}\right)^9} 1.01^{0.8} - 1.2 \frac{1-\left(\frac{p}{1-p}\right)}{1-\left(\frac{p}{1-p}\right)^9} 7.99^{0.8} > 0$ in this case.

accumulate gains of $x = 1.01$ or losses of -7.99. Observe that the gains and losses include the transfer $\tau^\star$. Since the gambler can make a positive expected payoff (the above strategy need not be optimal), he will accept the casino's inducement $\tau^\star$ and start gambling. From the casino's perspective, the inducement is profitable as even if the gambler plays only once, the casino's expected profit is strictly positive (note that the expected value of each lottery for the casino is $0.027 > \tau^\star$).

In the interaction that is described above, the gambler is allowed to continue gambling after the first lottery. If instead the gambler commits to participate in *exactly one lottery*, then he does not find the transfer beneficial. In fact, there exists no transfer τ' such that both the gambler is willing to accept τ' in return for his participation in exactly one lottery, and the casino is willing to pay τ' for the gambler's participation. Such a transfer is beneficial for the casino only if $\tau' \leq 1 - 2p$. The gambler accepts such a transfer only if

$$p\left(1+\tau'\right)^{\alpha} - (1-p)\,\lambda\left(1-\tau'\right)^{\beta} \geq 0 \tag{4}$$

The LHS of (4) is smaller than the LHS of (5) if $\tau' \leq 1 - 2p$.

$$p\,(2-2p)^{\alpha} - (1-p)\,\lambda\,(2p)^{\alpha} \geq 0 \tag{5}$$

But (5) cannot hold for $p < 0.5$. Hence, any transfer that incentivizes the gambler to participate will not be offered by the casino.

The reason that the gambler's ability to continue gambling (after participating in the first lottery) made $\tau^\star$ beneficial to both parties is related to the "break-even" effect. The gambler's diminishing sensitivity to losses means that he values the ability to break even in case he loses in the first lottery (i.e., $V(a^\star, -1, p, \delta) > U(-1)$). Thus, the ability to break even lowers the necessary compensation that is required to make the gambler play. Moreover, the break-even effect increases the casino's willingness to pay for participation as it increases the expected number of unfair lotteries that the gambler participates in (e.g., if the gambler loses the first lottery, then he will play again in an attempt to break even).

4 Dynamic Inconsistency

Our analysis in the previous sections assumed that the gambler is dynamically consistent (i.e., his reference wealth does not change over time). In this section, we shall consider the possibility that, occasionally, the gambler internalizes his profits and updates his reference wealth accordingly. We capture this idea by assuming that, in each period t, there is a probability π that the reference point is exogenously updated to the gambler's current wealth.[7] Thus, the update affects $x_t = w_t - r_t$ and, potentially, the gambler's behavior.

Let us clarify the timeline within a period. The gambler starts period t with wealth w_t and reference point r_{t-1}. With probability π (respectively, $1-\pi$), he updates his reference point to $r_t = w_t$ (respectively, $r_t = r_{t-1}$). He then chooses whether or not to participate in a lottery. After the lottery is realized, w_t is updated to w_{t+1}. We assume that once w_{t+1} is realized, if the gambler does not want to participate in additional lotteries, then he leaves the casino (i.e., ends the game) and does not wait for the next update to his reference point.[8]

In order to analyze the effects of the changes in the gambler's reference wealth on his behavior, we apply a multi-selves approach (Strotz, 1956) and model the interaction as a game played among different *selves*. Each reference wealth update induces a new self that makes the decisions (whether to stop or continue gambling) on behalf of the gambler until the next update. Each self cares about the gains with respect to his own reference wealth (i.e., the gambler's wealth at the time that self started to play). Observe that different selves may have the same wealth but a different reference wealth and, therefore, different gains with respect to their reference wealth. This may lead to dynamically inconsistent behavior as the selves' preferences and decisions depend on their gains rather than on their absolute wealth.

We split the analysis into two parts as we consider both the case of a gambler who is unaware of the possibility that his reference point will be updated, and the case of a gambler who is aware of these changes. We shall refer to the former type of gambler as a *naive* gambler and to the latter type as a *sophisticated* one. In the case of a naive gambler, each self plays the game as if he were the last self who will be called to play. That is, each self's optimal behavior is identical to the dynamically consistent (i.e., $\pi = 0$) gambler's optimal behavior. The sophisticated gambler's behavior is not so straightforward as each of the sophisticated gambler's "selves" plays the game taking into account his successors' behavior.[9]

Our main objective in this section is to compare the expected number of lotteries played by naive dynamically inconsistent gamblers with the expected number of lotteries played by sophisticated dynamically inconsistent gamblers.

[7] The idea of unpredicted changes to the reference point appears in a slightly different context in Barkan and Busemeyer (2003).

[8] If the gambler does not leave the casino, his reference point will eventually be updated to his current wealth and he may continue gambling.

[9] We do not restrict the sophisticated gambler to using stationary strategies as, in the present case, this restriction entails loss of generality.

The next result compares the dynamically consistent gambler's gambling plan with the dynamically inconsistent sophisticated gambler's selves' gambling plans (recall that the naive gambler, being unaware of his self-control problem, plans to play exactly like the dynamically consistent gambler). In the next result we focus on the generic case in which the dynamically consistent gambler's strategy is unique.[10] The result establishes that each of the sophisticated gambler's selves stops gambling before the dynamically consistent gambler does.

Proposition 3 *Fix arbitrary* $p < 0.5, \pi > 0$, *and* $\delta < 1$ *such that the dynamically consistent gambler's optimal strategy is unique: he stops gambling after accumulating gains of* $h \geq 0$ *or losses of* $-l \leq 0$. *Fix a subgame perfect Nash equilibrium of the multi-selves game and a sophisticated gambler's self* j. *If self* j *reaches gains levels of* $-l$ *or* h, *then he stops gambling and leaves the casino.*

Each of the sophisticated gambler's selves is aware of possibility that the reference wealth will be updated and a "new self" will be called to play. Effectively, this is a constraint on a self's ability to implement his preferred gambling plan as new selves may not stop gambling in instances in which "preceding selves" would want them to do so. This constraint incentivizes the sophisticated gambler to leave the casino before the dynamically consistent gambler does in order to prevent future selves from over-gambling. Leaving the casino serves as a "commitment device" in such cases.

We now illustrate another reason for the sophisticated gambler to leave the casino: the inability to break even. When the reference-wealth updates are frequent (e.g., when π is relatively close to 1) and succeeding selves do not gamble, a self is unlikely to receive an opportunity to break even if he loses in the first lottery. In such a case, the self that is playing prefers to leave the casino as participation in one lottery is never attractive to a loss-averse economic agent (since $u(x) < v(x)$).

The naive gambler's behavior is different. Each of his selves erroneously believes that the reference wealth will not be updated in the future and, therefore, plans to play exactly as the dynamically consistent gambler does (i.e., to use a left-skewed stopping rule). However, the naive gambler has *self-control* problems that may prevent him from implementing his preferred plan of action. Unlike the sophisticated gambler, he is unaware of these problems and, therefore, does not design his original plan to mitigate these problems.

The next result is a corollary of Proposition 3 (again, we focus on the generic case in which the dynamically consistent gambler's optimal strategy is unique).

[10]Observe that small changes in p break the gambler's indifference in instances in which he is indifferent between different gambling plans.

Proposition 4 compares the expected number of lotteries in which naive and sophisticated gamblers participate. It establishes that, in expectation, sophisticated gamblers participate in fewer lotteries than naive ones.

Proposition 4 *Fix p and δ such that the dynamically consistent gambler's optimal strategy is unique. For every $\pi > 0$, the expected number of lotteries in which the naive gambler participates is weakly smaller than the expected number of lotteries in which the sophisticated gambler participates.*

The intuition for this result is as follows. Leaving the casino is a commitment device as it guarantees that future selves will not be able to gamble. However, there is a cost to this commitment: leaving the casino too early prevents the gambler from implementing his preferred gambling plan. The sophisticated gambler, who predicts his "self-control" problem, leaves the casino earlier than the naive one who does not think that he needs such a commitment as he erroneously believes that his preferences are dynamically consistent.

We shall now compare the expected number of lotteries in which dynamically consistent gamblers participate with the expected number of lotteries in which dynamically inconsistent gamblers participate. Unlike in the previous comparison (i.e., sophisticated vs. naive), there is no clear-cut answer. The fact that dynamically inconsistent players may play more than dynamically consistent players is quite standard (see, e.g., Ebert and Strack, 2015). However, in our model, dynamically consistent players may participate, in expectation, in fewer lotteries than dynamically inconsistent *naive* players. We illustrate this effect in the next example. Observe that, by Proposition 4, in such cases, in expectation, the sophisticated dynamically inconsistent gambler plays in fewer lotteries than the dynamically consistent one as well.

Example: Under-gambling by naive gamblers

Consider the preferences given in (1), set $\lambda = 1.1, \alpha = \beta = 0.87, p = 0.49, \pi = 1$, and consider the $\delta = 1$ limit. The optimal stopping rule a for a dynamically consistent gambler must maximize $lim_{\delta \to 1} V\left(0, a, p, \delta\right)$=

$$\frac{1-\left(\frac{1-p}{p}\right)^{l}}{1-\left(\frac{1-p}{p}\right)^{h+l}} h^{0.87} - 1.1 \frac{1-\left(\frac{p}{1-p}\right)^{h}}{1-\left(\frac{p}{1-p}\right)^{h+l}} l^{0.87}$$

It is possible to show that the optimal stopping rule is to stop after accumulating gains of $h = 1$ or losses of $-l = -7$ with respect to the gambler's reference wealth.

In expectation, the number of lotteries in which the gambler participates is (the calculation is a simple application of the well-known gambler's ruin problem)

$$\frac{l}{1-2p} - \frac{l+h}{1-2p} \frac{\left(\frac{1-p}{p}\right)^{l} - 1}{\left(\frac{1-p}{p}\right)^{l+h} - 1} = 7.27$$

Since the naive gambler believes that his reference point will never be updated, he tries to implement the dynamically consistent gambler's optimal plan (i.e., stopping after accumulating gains of $x = 1$ or losses of $x = -7$). In each period, the naive gambler internalizes his profit (i.e., his reference wealth becomes his present wealth). Therefore, he stops gambling after the first lottery in which he wins (he then reaches gains of $x = 1$ relative to his reference wealth). Hence, the expected number of lotteries he participates in is $\sum_{z=1}^{\infty} p\,(1-p)^{z-1}\, z = 2.04$.

In expectation, the naive gambler participates in fewer lotteries then the dynamically consistent gambler. The reason for this effect is that the "upper bound" of the gambler's strategy is at $x = 1$. Thus, whenever the reference point is updated, the gambler goes back to the starting point (i.e., x=0) and gets closer to hitting that bound and stopping. Observe that if the naive gambler's optimal strategy were to stop after accumulating gains of $x > 1$, then he would *never stop gambling* as he would never accumulate such gains with respect to his constantly changing reference wealth.

5 Related Literature

Loss aversion is one of the most established departures from classic expected utility theory. It was introduced by Kahneman and Tversky (1979) and was applied to explain well-documented phenomena such as the endowment effect (Thaler, 1980), the equity premium puzzle (Benartzi and Thaler, 1995), and low price variance among differentiated products (Heidhues and Kőszegi, 2008).[11]

Barberis (2012) identified that probability distortion à la cumulative prospect theory (Kahneman and Tversky, 1992) creates a taste for right-skewed lotteries. He showed that it is possible to generate such a skewed lottery from a finite sequence of 50:50 binary lotteries and that probability distortion leads to dynamic inconsistency when a gambler faces a sequence of such lotteries. The inconsistency follows from the fact that, initially, the gambler puts different weights on different final outcomes (e.g., he puts a relatively high weight on highly unlikely

[11] Other prominent applications of loss aversion (in different contexts) appear in Herweg and Schmidt (2014), Karle and Peitz (2014), Carbajal and Ely (2016), and Rosato (2016).

events such as large profits). As time progresses, the likelihood of different events changes and so do the relative weights that the gambler assigns to these events.

Barberis's (2012) dynamic inconsistency is the point of departure of Ebert and Strack (2015, 2016) whose results are related to the analysis in Section 4 of the present paper. Ebert and Strack (2015) study a slightly different setting from Barberis's and obtain a surprising result: under mild assumptions on the probability distortion, a naive[12] gambler never stops gambling.[13] Ebert and Strack (2016) study the behavior of a sophisticated player and obtain another striking result: the only strategy that a sophisticated gambler can execute is to never gamble. The gambler in Ebert and Strack (2016) underweights highly likely events and this prevents him from executing any strategy that involves gambling until he accumulates profits of $h > 0$. Such strategies are non-executable since once the gambler starts winning, h becomes more likely and, therefore, underweighted so that the gambler prefers to stop immediately.

6 Concluding remarks

This paper presented a model in which a loss-averse player decides when to stop an infinite sequence of unfair lotteries. We showed that the optimal gambling plan for a loss-averse player is a left-skewed stopping rule (i.e., a rule that guarantees a small prize with a relatively high winning probability), and that it is always possible to find a sequence of unfair lotteries that a loss-averse player would be willing to participate in.

We established that stochastic updates to the reference wealth lead to dynamically inconsistent gambling behavior. Since the player's preferences are defined over gains and losses with respect to his reference wealth, any change in his reference point affects his behavior. Knowing that they will gamble too much in the future, sophisticated players, who are aware of this "self-control" problem, try to mitigate it by planning to leave the casino earlier than dynamically consistent players and earlier than naive players, who are unaware of their self-control problem.

Left-skewed stopping rules often induce left-skewed lotteries (for $p < 0.5$, a left-skewed stopping rule may induce a right-skewed lottery). This finding allows us to distinguish between the two main components of prospect theory: loss aversion and probability distortion. While it is known that overweighting of

[12]In these papers, a gambler is said to be naive (respectively, sophisticated) if he is unaware (respectively, aware) of his dynamic inconsistency.

[13]Henderson et al. (2017) show that this is no longer the unique prediction when the gambler is allowed to use randomized strategies.

unlikely events and underweighting of highly likely events imply preferences for right-skewed lotteries, loss aversion implies a taste for left-skewed lotteries.

References

[1] American Gaming Association (2017): "State of the States 2017," retrieved from https://www.americangaming.org/research/reports/state-states-2017-aga-survey-casino-industry.

[2] Barberis, N. (2012): "A Model of Casino Gambling," *Management Science,* 58, 35–51.

[3] Barkan, R. and Busemeyer, J. (2016): "Modeling Dynamic Inconsistency with a Changing Reference Point," *Journal of Behavioral Decision Making,* 16, 235–255.

[4] Benartzi, S. and Thaler, R. (1995): "Myopic Loss Aversion and the Equity Premium Puzzle," *Quarterly Journal of Economics,* 110, 73–92.

[5] Blackwell, D. (1965): "Discounted Dynamic Programming," *The Annals of Mathematical Statistics,* 36, 226–235.

[6] Carbajal, J. and Ely, J. (2016): "A Model of Price Discrimination under Loss Aversion and State-Contingent Reference Points," *Theoretical Economics,* 11, 455–485.

[7] Conlisk, J. (1993): "The Utility of Gambling," *Journal of Risk and Uncertainty,* 6, 255–275.

[8] Ebert, S. and Strack, P. (2015): "Until the Bitter End: On Prospect Theory in a Dynamic Context," *American Economic Review,* 105, 1618–1633.

[9] Ebert, S. and Strack, P. (2016): "Never, Ever Getting Started: On Prospect Theory without Commitment," *SSRN working paper.*

[10] Friedman, M. and Savage, L. J. (1948): "Utility Analysis of Choices Involving Risk," *Journal of Political Economy,* 56, 279–304.

[11] Grinstead, C. and Snell, J. (1997): *Introduction to Probability,* American Mathematical Society.

[12] Heidhues, P. and Kőszegi, B. (2008): "Competition and Price Variation When Consumers Are Loss Averse," *American Economic Review,* 98, 1245–1268.

[13] Henderson, V., Hobson, B., and Tse, A. (2017): "Randomized Strategies and Prospect Theory in a Dynamic Context," *Journal of Economic Theory,* 168, 287–300.

[14] Herweg, F. and Schmidt, K. (2014): "Loss Aversion and Inefficient Renegotiation," *Review of Economic Studies,* 1, 1–38.

[15] Kahneman, D. and Tversky, A. (1979): "Prospect Theory: An Analysis of Decision under Risk," *Econometrica,* 47, 263–292.

[16] Kahneman, D. and Tversky, A. (1992): "Advances in Prospect Theory: Cumulative Representation of Uncertainty," *Journal of Risk and Uncertainty,* 5, 297–323.

[17] Karle, H. and Peitz, M. (2014): "Competition under Consumer Loss Aversion," *Rand Journal of Economics,* 45, 1–31.

[18] Markowitz, H. (1952): "The Utility of Wealth," *Journal of Political Economy,* 60, 151–158.

[19] Rosato, A. (2016): "Selling Substitute Goods to Loss-Averse Consumers: Limited Availability, Bargains and Rip-offs," *Rand Journal of Economics,* 47, 709–733.

[20] Strotz, R. H. (1956): "Myopia and Inconsistency in Dynamic Utility Maximization," *Review of Economic Studies,* 23, 165–180.

[21] Thaler, R. (1980): "Toward a Positive Theory of Consumer Choice," *Journal of Economic Behavior and Organization,* 1, 39–60.

Appendix

Proof of Lemma 1

First, assume by negation that never stopping (denoted by a') is an optimal strategy for the gambler. This implies that $V(a', x, p, \delta) \geq U(x)$ for every gains level x. Fixing a', increasing p would only increase the gambler's value. Therefore, $V(a', x, 0.5, \delta) \geq U(x)$ for every x. Hence, $V(a', 0, 0.5, \delta) \geq 0$. However,

$$V\left(a', 0, 0.5, \delta\right) = (1-\delta)\,[\left(\frac{1}{2}u(1) - \frac{1}{2}v(1)\right) + \frac{1}{2}\delta\left(\frac{1}{2}u(2) - \frac{1}{2}v(2)\right) + \frac{1}{4}\delta^2\left(\frac{1}{2}u(3) - \frac{1}{2}v(3)\right) + \frac{3}{4}\delta^2\left(\frac{1}{2}u(1) - \frac{1}{2}v(1)\right) + ...] \quad (6)$$

Since $V(a', 0, 0.5, \delta)$ is a combination of "50-50" fair lotteries and the gambler is loss averse (i.e., $u(x) < v(x)$), it follows that $V(a', 0, 0.5, \delta) < 0$. This is in contradiction to a' being optimal.

Second, assume that the gambler stops gambling once he accumulates gains of $h > 0$ and only then. Denote this strategy by a^h and assume that it is optimal.

Since the gambler is free to stop gambling whenever he wishes, if a^h is optimal, the value $V\left(a^h, x, p, \delta\right)$ is weakly increasing in δ. Therefore,

$$lim_{\delta\to 1} V\left(a^h, 0, p, \delta\right) \geq V\left(a^h, 0, p, \delta'\right) \geq U(0) = 0$$

for every $\delta' < 1$. At the $\delta = 1$ limit, the gambler's problem is an application of the well-known "gambler's ruin" problem (for a textbook treatment, see, e.g., Grinstead and Snell, 1997). If $p < 0.5$, with a strictly positive probability $1 - \left(\frac{p}{1-p}\right)^h$ the gambler becomes infinitely poor, and since $v(x)$ is unbounded, $Lim_{\delta\to 1} V(a', 0, p, \delta) < 0$. Therefore, it is not optimal to gamble at $x = 0$, which is in contradiction to the optimality of a^h.

Finally, assume by way of negation that stopping after accumulating losses of $-l$ and only then is optimal for the gambler. Denote this strategy by a^l. Again, consider the $\delta = 1$ limit and note that

$$lim_{\delta\to 1} V\left(a^l, 0, p, \delta\right) \geq V\left(a^l, 0, p, \delta'\right) \geq U(0) = 0$$

for every $\delta' < 1$. Since $p < 0.5$, at the $\delta = 1$ limit, with probability 1, the gambler is "ruined" at the end of the game under a^l. That is, $lim_{\delta\to 1} V\left(a^l, 0, p, \delta\right) = U(-l) < 0 = U(0)$. This is in contradiction to the optimality of the gambler's strategy a^l.

In conclusion, if a is optimal, then there must be a gains level $h \geq 0$ and a losses level $-l \leq 0$ such that, under a, the gambler stops once he accumulates gains of h or losses of $-l$.

Proof of Proposition 1

By Lemma 1, we can focus on stopping rules. We shall prove Proposition 1 by showing that right-skewed stopping rules and symmetric stopping rules (i.e., stopping after accumulating gains of h or losses of $-l$, where $h \geq l > 0$) cannot be optimal. Assume by negation that the gambler stops after he accumulates gains of h or losses of $-l$. We will show that if $h \geq l > 0$, then the gambler would rather stop gambling once he reaches his reference wealth (i.e., he would prefer not to gamble at all).

Denote an optimal stopping rule by $a^\star$ and recall that $V(a^\star, 0, p, \delta) \geq 0$. Increasing p and δ would only increase $V(a^\star, 0, p, \delta)$. Therefore, $lim_{\delta\to 1} V(a^\star, 0, 0.5, \delta) \geq 0$. The latter expression is given in the LHS of (7):

$$\frac{l * u(h)}{h + l} - \frac{h * v(l)}{h + l} \geq 0 \tag{7}$$

Rearranging,

$$\frac{u\left(h\right)}{h}\geq\frac{v\left(l\right)}{l} \tag{8}$$

By loss aversion, inequality (8) cannot hold for $h=l>0$ as $\frac{v(l)}{l}>\frac{u(l)}{l}$. Since $\frac{v(x)}{x}$ is weakly decreasing in x and $u\left(x\right)<v\left(x\right)$ for all x, inequality (8) cannot hold for $0<l<h$.

Proof of Proposition 2

Since the gambler is dynamically consistent, it is sufficient to show that there exists a strategy $a^{\star}$ such that $lim_{(p,\delta)\rightarrow(0.5,1)}V\left(a^{\star},0,p,\delta\right)>U\left(0\right)$. Let $a^{\star}$ be a stopping rule under which the gambler stops after accumulating gains of $h>0$ or losses of $-l<0$. This condition is given by

$$lim_{(p,\delta)\rightarrow(0.5,1)}V\left(a^{\star},0,p,\delta\right)=\frac{l*u\left(h\right)}{h+l}-\frac{h*v\left(l\right)}{h+l}>0 \tag{9}$$

It is possible to rearrange (9) and to obtain $\frac{u(h)}{h}>\frac{v(l)}{l}$. Since $lim_{x\rightarrow\infty}\frac{v(x)}{x}=0$, it is always possible to find l and $h<l$ such that (9) holds and the gambler prefers stopping at gains of h and losses of l to never playing.

Proof of Proposition 3

In order to prove this result, we shall show that if the dynamically consistent gambler stops gambling at gains (or losses) of x, then self j must stop gambling at that gains level as well. Denote the dynamically consistent gambler's optimal strategy by $a^{\star}$. If the dynamically consistent gambler stops gambling at gains (or losses) of x according to $a^{\star}$, then $V\left(a^{\star},x,p,\delta\right)=U\left(x\right)$. By the assumption that the dynamically consistent gambler's optimal strategy is unique, if the strategy a' includes gambling at gains of x, then $V\left(a',x,p,\delta\right)<V\left(a^{\star},x,p,\delta\right)=U\left(x\right)$.

Consider an arbitrary self j and fix an arbitrary profile $\left(a_k\right)_{k\neq j}$ of strategies played by the other selves. Consider a gains level x and assume that, given some history, self j does not leave the game upon reaching gains level x. In a subgame perfect Nash equilibrium, self j's strategy is a best response to $\left(a_k\right)_{k\neq j}$. Hence, since he can always leave the casino with gains of x, self j expects to obtain a utility of at least $U\left(x\right)$ conditional on reaching x. Since the dynamically consistent gambler can imitate the behavior of self j's successors, it must be that $V\left(a',x,p,\delta\right)\geq U\left(x\right)$ for some strategy α' in which the dynamically consistent gambler gambles at gains level x. This leads to a contradiction as $V\left(a',x,p,\delta\right)<$

$V(a^\star, x, p, \delta) = U(x)$ for every strategy $a' \neq a^\star$ that includes gambling at gains of x.

Proof of Proposition 4

Since each of the naive gambler's selves believes that he is the last self to play, their behavior is identical to the dynamically consistent gambler's behavior. By Proposition 3, at any gains level at which a naive gambler's self stops gambling and leaves the casino, a sophisticated gambler's self leaves the casino as well.