Quantitative Marketing and Economics
https://doi.org/10.1007/s11129-025-09293-8

CRM Targeting with reference-dependent sensitivities: Evidence from the casino industry

Wayne Taylor[1] · **Jonathan Zhang**[2]

Received: 20 February 2024 / Accepted: 2 January 2025

Abstract
This research explores heterogeneity in customers' reference-dependent sensitivities using rich, individual level CRM data from a large casino in the U.S. and discusses implications for targeting decisions. We use a unique panel dataset of over 12,000 slot machine gamblers over 14 years and model heterogeneity in reference-dependent sensitivities at the individual level using a hierarchical Bayesian model. This analysis focuses on gains and losses relative to three reference points unique to the casino industry but conceptually extends to many other settings such as the financial services industry and hospital- ity: gambling outcomes relative to 1) zero, 2) prior trip outcomes, and 3) expected losses based on the house advantage of the slot machines. Firms can use heterogeneous reference-dependent sensitivities to improve their targeting decisions by considering the sequences of gambler outcomes in tandem with gamblers' individual sensitivities to marketing promotions. In our empirical application, we estimate that incorporating individual-level reference-dependent sensitivities improves targeted offer profitability by at least 19.8% relative to a comparable RFM model, depending on the offer type.

Keywords Casino · Customer relationship management · Targeting · Reference-dependent sensitivities

✉ Wayne Taylor
wjtaylor@smu.edu

Jonathan Zhang
jonathan.zhang@colostate.edu

1 Southern Methodist University, Dallas, USA

2 Colorado State University, Fort Collins, USA

1 Introduction

This paper focuses on the marketing problem of whom to target in settings where customers exhibit reference- dependent sensitivities to their interactions with the firm. This situation is most common in industries where there is naturally occurring variation in the customer experience over time, which can shape customer outcomes relative to these reference points and influence future decisions. For example, in the financial services industry a customer's satisfaction may be correlated with quarterly gains and losses of their portfolio. Their satisfaction may depend on which reference point the customer compares performance to: the prior quarter's performance, overall market performance, their own expectations, or some combination of all three. Each customer may be sensitive to different reference points that informs their judgment of the firm and their future engagement. Furthermore, at each reference point individual sensitivities may vary.

In service-based industries especially, variation in the customer experience over time is more pronounced. For example, rideshare services, hotels and restaurants, and grocery delivery tend to exhibit higher variation in customer experiences compared to firms selling packaged goods. By analyzing both the series of outcomes relative to different reference points and the customer responses to these outcomes, firms may be able to improve their targeting decisions.

In this paper we show that reference-dependent sensitivities explain a substantial portion of next-trip casino gambler play volume and demonstrate the value of incorporating reference-dependent sensitivities into the casino's targeting decisions. In our dataset, each gambler experiences a net gain or loss at the end of their trip to the casino. We quantify these outcomes as gains or losses relative to three distinct reference points that are easily measured by the casino, appropriate for the empirical setting, and motivated by behavioral research. The first reference point is zero – how much did the gambler win or lose on this trip? The second reference point is the prior trip outcome – what was the gain or loss relative to the last trip outcome? The third and final reference point is the expected loss of the gambler based on the slot machine they played – that is, as all slot machines are programmed to take a percentage of a gambler's playing volume, did the gambler get "lucky" by losing a relatively low amount given the volume played? We estimate how gains and losses relative to these three reference points influence a gambler's play volume on their return trip. We use these estimates to suggest improvements to a casino's targeting decisions.

Our work contributes to the broad empirical literature on customer relationship management (CRM) and targeted marketing actions (Deng & Mela, 2018; Dong et al., 2009; Dubé & Misra, 2023; Kumar et al., 2011; Simester et al., 2006; Smith et al., 2023; Tuchman et al., 2018; Zantedeschi et al., 2017; Zhang et al., 2014). Customized marketing has been a cornerstone in marketing for decades (Rossi et al., 1996), where firms identify and target distinct segments of customers in order to improve their profitability relative to mass marketing strategies. Traditionally, firms segment customers on first order transaction variables such as purchase quantities, recency, frequency, and spend. More recently, researchers have explored additional metrics that can offer greater insight into customer preferences and behaviors to refine targeting decisions. Taylor and Hollenbeck (2021) use the locations of competitor stores to improve the targeting of a loyalty program. Padilla et al. (2023) explores how customer footprints from customer journeys can inform retention and engagement strategies. Dew et al. (2023) examine customer routines to aid targeting decisions. These studies underscore the growing use of alternative measures to improve targeting decisions.

We continue this trend by considering customers' reference-dependent measures inspired by past behav- ioral research. Whether or not our reference-dependent measures correlate with future gambler behavior is an open empirical question which we test using data from a large U.S. casino. The casino industry provides a fitting context for studying individual reference-dependent sensitivities, as it is characterized by repeated and randomly varying experiences of gains and losses that can be easily quantified. While the concept of reference-dependence applies to many industries, contributions to the casino industry alone are significant given the industry's size and economic impact.

The primary objective of casino marketing is to induce incremental play volume on return trips, condi- tional on a return visit.[1] Therefore, we estimate changes in play volume as a function of gains and losses relative to our three reference points previously described. We augment this specification to allow marketing to moderate the sensitivities to these reference points.[2] Specifically, we use trip level gambling outcomes from 12,150 slot machine players over 14 years, focusing on the gamblers that return to the casino on a recurring basis.[3] We start with a homogeneous model to show that including gains and losses relative to the three reference points adds considerable fit above standard recency-frequency-monetary (RFM) variables currently used by a typical casino. Next, we employ a hierarchical Bayesian model to allow for individual-level hetero- geneity in reference-dependent sensitivity which provides granular behavioral insight. We show that gains in reference-dependent targeting decisions are economically substantial relative to these benchmark models and the casino industry's current practice.

The contributions of our study are threefold. First, we provide empirical evidence of reference-dependent sensitivity using actual transaction data over a long period of time. Our results provide insight into both the relevance and heterogeneity of reference-dependent sensitivity in an externally valid setting. The large-scale field context not only complements related work in psychology with their experiment-based and survey- based approaches (Brown et al., 2024; Mrkva et al., 2020),

[1] This insight is based on discussions with multiple industry executives. Similar to many discretionary services sectors such as travel and vacation, customers' return decisions are often determined by many other factors outside of the firm's control, such as work schedules and vacation planning.

[2] While it is tempting to suggest casinos model *within-trip* play behavior in an attempt to deliver offers in real time as opposed to *across-trip* activity, given the sophistication of most casinos this is simply not externally valid nor as potentially useful in terms of within-trip data analytics and marketing intervention.

[3] That is, our analysis focuses on how outcomes influence future play volume, not the decision of whether or not to return.

it contributes to work in economics investigating endogenous reference point construction (Kőszegi & Rabin, 2006), how reference points affect consumption and risk attitudes (Kőszegi & Rabin, 2007), and how experiences affect inter-temporal consumer dynamics (Hendel & Nevo, 2013).

Our second contribution is a byproduct of the first and grounded in improving CRM practice: if gains and losses relative to specific reference points influence gambler play volume, casino marketers should be able to improve their targeting strategies by accounting for gamblers' sensitivities to these outcomes. In our empirical application, we estimate that targeting profitability improves by at least 19.8% relative to an individual-level RFM model with no reference-dependent variables. Intuitively, certain sequences of wins and losses may reduce the need for costly marketing offers intended to increase play volume. Understanding how gamblers react to outcomes informs the casino of which offers are more or less effective at driving incremental return trip play volume. As previously noted, incorporating reference-dependent sensitivities into targeting decisions easily extends beyond the casino context, especially in many service-based industries where the customer experience varies naturally due to factors outside of the firm's direct control.

Finally, and more broadly, this paper advances our understanding of the interplay of diverse elements that contribute to customers' overall experiences and subsequent decision-making. These approaches al- low firms to gain deeper insights into their customers, track the relevant experience metrics, and develop performance-enhancing marketing initiatives based on these experiences. Our study responds to Lemon and Verhoef (2016)'s request for further research on the "conceptualization, drivers, and consequences of customer experience" as well as "customer experience measurement" (e.g., "How can CX [customer experience] be assessed while considering its rich, multidimensional nature?", pg 87).

1.1 The casino industry: Institutional background, "Theoretical Win" (Theo), and current CRM practice

Our data comes from a large U.S. casino. The gaming industry is a substantial component of the U.S. and global economy. Casino gaming revenues have reached a new high, exceeding \$75 billion in 2022 in the U.S and \$260 billion globally (American Gaming Association, 2022). With nearly 1,000 casinos operating in 44 states, Americans spend more money on casino entertainment than they do on spectator sports such as football, baseball, basketball, and soccer. A 2022 survey of more than 1,000 Americans show that 84% have been to a casino. 52% have gone to a casino more than 6 times, and 17% have been more than 30 times in their lives (PlayUSA, 2022).

Despite the size of the gaming industry, little research has examined improving its CRM practice. Narayanan and Manchanda (2012) use a hierarchical Bayesian model to identify gamblers who exhibit irra- tional behaviors and measured how they respond to marketing. The closest paper to this analysis is Nair et al. (2017), who provide a discrete choice model of return, play volume, and response to promotions with the goal of optimizing segmentation and targeting. These papers, however, do not consider individual-gambler reference points, discuss how reference-point sensitivities could affect return trip behavior, or describe the interactions between these sensitivities and marketing offers. Unlike Nair et al. (2017), we forego modeling the decision to return – in our data we did not find any meaningful relationship between reference-dependent sensitivities, marketing offer values, and return trip timing, so we exclude this in the interest of brevity. As we detail below, the gaming industry offers an ideal context to study reference-point sensitivities because customers' experience outcomes are exogenous: while the expected value in the long run is negative on any single trip outcomes vary substantially. Consequently, there is a need for increased attention to casino marketing to better understand the potentially unique behaviors of outcome experiences in reference formation and develop more effective marketing resource allocation strategies.

In the casino industry, the concept of "theoretical win" or "theo" is the primary metric used to determine a customer's value.[4] Because of the chance nature of casino games, theo represents the amount a casino expects to win from a player based on their activity and the built-in house advantage for each game. It is a proxy measure of customers' spending and engagement. Casinos use the following to determine a customer's theo: Theo = Average Bet x House Advantage x Time Played x Speed of Play.

The average bet is the average wager placed by the gambler, for example the bet per spin at a slot machine. The house advantage is the built-in statistical edge that the casino has for each game. Time played is the duration of the gambling session, and the speed of play represents the number of bets placed per unit of time. For example, if a player spends 2 hours at a slot machine that has a house advantage of 6% and bets \$1 per spin at a speed of one bet per minute, the theoretical win for the casino from that session of play is: \$1 average bet $\times$ 6% house advantage $\times$ 2 hours played $\times$ 60 wagers/hour = \$7.20 in theo.

For casinos' CRM practice, theo is used to segment gamblers and represents their worth in expectation. Gamblers with higher theo are deemed more valuable, as they are anticipated to generate greater revenue for the casino in the future. Because of this, marketing offer values are typically based on theo. This prevailing practice was confirmed with interviews with two casino executives, who do not target customers based on actual wins and losses – "We only segment by expected loss [of the customer]". When asked why not use actual outcomes, one manager mentioned that using theo is the industry norm and that "Using actual outcomes never really crossed our minds." Another executive alluded to the current practice's simplicity and revealed that their CRM approaches are limited by the skills of their analysts – "Our analysts are not sophisticated enough to do this." On the few instances when the industry integrates actual outcomes, it is not accomplished with the necessary granularity or through a statistical approach which would facilitate systematic decision support. One executive shared – "Sometimes we use average actual loss combined with theo, but it is an average over multiple trips – never individual trip

[4] Within the industry, "theo" may also be interchangeably referred to as "theoretical loss" (from the player's perspective), "theoretical win" (from the casino's perspective), "play volume", "theo loss", and "theo win". In all cases it is always a strictly positive value, regardless of terminology.

loss." They confirmed that except for a few instances involving significant individual losses, casinos generally do not take players' actual winning or losing outcomes into account when making marketing decisions.[5] While the casino executives we interviewed recognized the informational value of players' actual gambling outcomes, they often expressed surprise at the thought of incorporating these outcomes systematically into their marketing processes. The insights from our interviews highlight a potential opportunity to improve CRM in the casino industry through reference- dependent targeting. Of course, some casinos may already deploy more sophisticated schemes than what our interviewees use or what we propose (such as using additional metrics of engagement).[6] However, given that the executives we interviewed are each from a large and established destination casino we expect their views to reflect the norm rather than the exception.

The primary reason for this hesitation is because theo, understandably, eliminates the randomness of gambling outcomes and better reflects potential spending and engagement. The logic from the casino's perspective is that "on average" offers will be appropriate for the gamblers. While it may be true that in aggregate actual losses will be nearly equivalent to theo, the fact remains that gamblers observe and react to their own individual outcomes. Decades of research in consumer behavior and psychology have consistently shown that gains and losses can profoundly impact customers' current experiences and shape their subsequent actions. As a result, the industry may be allocating resources sub-optimally by not properly accounting for such crucial information pertaining to the customer experience.

Currently, at a typical casino targeting depends on trip-level outcomes rather than individual bet activity. Even though a casino has access to gambling activity at the individual bet level, it is unlikely to help the average operator in a substantive way. First, there is the obvious concern of hyper-targeting, which might dissuade a casino from attempting to target within a given trip. Second, given the current sophistication level related to targeting based on the managers we spoke with, this type of micro-targeting is simply not a realistic suggestion at this point. These industry insights serve as motivators to our trip-level approach to model specification and estimation.

With this institutional background in mind, we propose that by understanding customers' reference- dependent sensitivity at the individual-level and examining interactions with marketing responses, casinos can enhance and refine their CRM practices, thus optimizing their existing marketing spend. In our empirical study, we demonstrate that integrating players' theo with their actual experiences and specifying outcomes in a reference-dependent manner can lead to more effective targeting decisions. Furthermore, with the backdrop of consumer privacy concerns (Bleier et al., 2020) we show that using relatively lean customer information, namely trip-level theo and actual outcomes could be an effective yet unobtrusive way to target gamblers.

---

[5] Discretionary incorporation of significant losses tends to be limited to high-roller table games rather than slot machines, since a "quick large loss" on a slot machine is far less common.

[6] For example, the field experiment in Nair et al. (2017)

2 Data

Our data comes from a large, destination casino in the United States. We have trip-level information including gambling outcomes and marketing offer redemption from 49,082 slot gamblers across over 14 years up to 2013. Gamblers tend to play either slot machines or table games exclusively (a relatively small number of gamblers mix game types), and in this analysis we focus on slot gamblers since the accounting at the individual level is exact to the penny, unlike table games.[7] We limit our analysis to the 12,150 gamblers with at least seven trips at the casino, which is the average number of visits to the casino across all slot players. We do this for a few reasons: first, while these observations represent only about 25% of all slot gamblers, they generate about 55% of all casino trips and are of greater interest to the marketing department. Second, we wanted to ensure a reasonable number of observations per gambler in order to establish reference points and enable individual level estimates. Finally, in this analysis we are not interested in *whether* gamblers return, but instead how outcomes influence future play for returning gamblers. While there may be dynamics that can influence return decisions, as mentioned earlier the decision of when exactly to return to a destination casino is often outside of the control of the firm and is a separate question out of scope of this research. For these reasons we focus on the set of gamblers where survival bias would be less likely to influence our results.[8]

Table 1 provides additional information on the gamblers of interest. The average trip length is about two and a half days, and gamblers play for about seven hours per trip with an average trip loss of about $500. The house advantage of the slot machine dictates the theo of the gambler and is simply the product of the house advantage multiplied with the time played, speed of play, and average bet. The ratio of player loss to theoretical loss across all trips is close to one on average (as expected) but the standard deviation is very large. This reaffirms that most trips are not long enough for actual losses to converge to expected losses. Theo will serve as our primary dependent variable of interest and is the primary indicator of gambler engagement at casinos.

Marketing plays a large role in casino CRM strategies. The two common types of offers sent to gamblers are either promotional credits (also called “promo credits”) or complimentary rooms (“room comps”). Promo credits represents free slot play: the promo credit has no cash value, but any winnings earned from using the promo credits can be retained by the gambler. From the casino’s perspective, theadvantage

[7] In table games, a pit boss approximates a gambler’s theoretical loss based on average bet, time played, and the house advantage of the game. While the comparisons between table game and slot gamblers could provide an interesting to the analysis, it is out of scope for this paper.

[8] Initially, we analyzed changes in the *timing* of the return trip, but these results were much noisier and fit the data poorly. This is not surprising given that for most gamblers at destination casinos trips tend to be planned in advance and revolve around work schedules or other life events. We corroborated our results on trip timing during our interviews with gamblers. Because of this, and in the interest of brevity, we did not include trip timing in our modeling framework and analysis but this could easily be included in situations where it might be more relevant.

Table 1 Summary Statistics

	Mean	SD
# of Trips (per gambler)	13.7	8.6
Age	59.5	10.6
Male	35.4%	n.a
Trip Statistics		
Trip length (days)	2.3	0.8
Hours played	6.7	5.2
Average bet	$1.58	$1.26
Player loss	$488	$416
Theoretical loss	$498	$408
Player Loss/Theoretical Loss	1.09	7.01
Promo Credits (43.4% of trips)		
Avg. value	$144	$181
Room Comps (37.2% of trips)		
Avg. value	$374	$156
Room nights	2.5	1.4

For trip statistics, trip information is first averaged by gambler, and then we present the averages and standard deviations across gamblers

of promo credits over complimentary rooms is that they are not subject to supply constraints (since there are only so many hotel rooms that can be comped). The disadvantage is that they are treated as a hard cost: whereas a $300 room comp may only cost the casino $30 in cleaning expenses, a $300 promo credit can cost the casino close to the full $300.[9] Combined promo credits and room comps on a single trip is not uncommon – this occurs on 62% of the trips where any offer is redeemed.

2.1 Marketing offers

Marketing offers tend to be a prominent aspect of casino CRM and are sent out frequently: as Table 1 shows, on more than a third of all trips gamblers redeemed either promotional credits or complimentary room nights. The value of the offer, and who qualifies for an offer, is determined by the "reinvestment rate", which reflects how much a casino is willing to spend on offers relative to theo. The reinvestment rate generally remains fixed over time and depends on the strategy set forth by casino management. For example, an 8% reinvestment rate means that for every $100 in theo, a gambler qualifies for $8 worth of offers. The casino will first decide to send, say, a $50 or $100 promo credit offer and then determine which gamblers qualify for an offer given their historical play using the last trip

[9] If the house advantage on a slot machine is 8%, when a gambler cycles $300 through the slot machine the casino will only keep $24 and the gambler can walk away with the rest as winnings. Of course, the gambler may continue to cycle the payouts back into the machine in which case the cost of the promo credit is reduced.

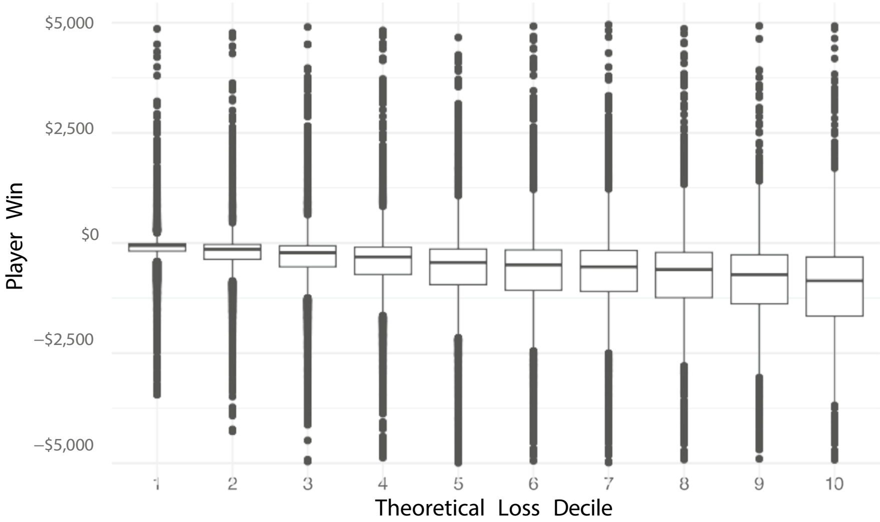

Fig. 1 Trip level player win by theoretical loss decile

theo.[10] Using theo to segment customers is the standard practice among casinos for offer generation. Since our data only contains offer redemptions, we do not directly observe the process of offer issuance. Since redemption requires an offer to have been received, we use redemption as a proxy for offer receipt, acknowledging that this is a noisy measure subject to random errors. The noise is driven by unobservable factors such as offer timing, gambler trip schedules, or the gambler's interest in specific types of offers.

This type of targeting strategy is ideal for our model estimation because offer values are based on last trip theoretical loss, not actual outcomes. The implication is that marketing offers are relatively exogenous with respect to our reference dependent heuristics, which are constructed based on actual player wins and losses. In addition, actual outcomes are only loosely correlated with theoretical loss. Figure 1 shows the distribution of trip-level player outcomes conditional on theoretical loss deciles. As theoretical loss increases, player win decreases (on average) – representing the effects of the house advantage. However, the box plots are very wide, indicating that conditional on a given theoretical loss there is considerable variation in actual, experienced outcomes across gamblers. It is not uncommon for a set of gamblers who experience a wide range of actual outcomes to receive similar offers because of similar theoretical losses. This unique combination of 1) targeting based on expected losses, not actual losses and 2) randomly generated outcomes from slot

[10] Based on discussions with the casino, last trip theo is a good approximation of the targeting rule. On rare occasions they may consider lifetime theo or some moving average but this is unusual. Because historical marketing offer information is not tracked well at the focal casino, we emphasize interpreting estimated marketing effects as approximate.

Table 2 Marketing offer redemption propensity and expected value

	Pr(Offer Redemption$_t$)		E[Offer Value$_t$]	
	Promo Credit	Room Comp	Promo Credit	Room Comp
Theo$_{t-1}$	0.0001***	0.0001***	0.0427***	0.0161***
	(0.000005)	(0.000005)	(0.0015)	(0.0020)
Casino Win$_{t-1}$	0.000005*	0.000002	0.0014	– 0.0016
	(0.000003)	(0.000003)	(0.0009)	(0.0012)
Trip Number	0.0213***	0.0102***	6.6273***	– 5.9383***
	(0.0011)	(0.0011)	(0.4594)	(0.6365)
Observations	60,750	60,750	22,905	21,352
Gambler FE	Yes	Yes	Yes	Yes

$^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

machines creates exogenous variation in marketing offers with respect to gamblers' realized gains and losses.

Even though offer values are exogenous with respect to actual outcomes, a remaining concern is whether the gamblers who redeem offers are different from gamblers who do not redeem offers. For instance, if we consider offer reception and offer redemption as:

$$\begin{aligned} \text{Receive Offer}_t &= \text{Theo}_{\text{t}-1} + \epsilon_t \\ \text{Redeem Offer}_{t+1} &= \text{Receive Offer}_t + \nu_t \end{aligned} \tag{1}$$

where t represents time since an offer has to be received before it can be redeemed on a trip. The error term ϵ_t reflects unobserved reasons why gamblers may not receive any offer conditional on theo; perhaps the casino started processing the offer list prior to the gambler's most recent trip or the casino simply decided to not send offers during a particular time period. After conditioning on last trip theo, we don't expect any systematic biases to influence offer reception. The error term ν_t are the unobserved reasons why some individuals redeem and others do not. Some gamblers may only travel during offer "blackout" dates (e.g., high demand dates), and are unable to use their offers. If the gamblers who redeem and substantially different from those who do not redeem, our estimates of marketing effectiveness could be biased.

We have empirical evidence to suggest that the gamblers who redeem are similar to those who do not redeem. First, for the vast majority of gamblers (82%), we observe both redemption and non-redemption trips – thus redeemers and non-redeemers are identical. The remaining 18% of gamblers are split evenly between gamblers who never redeem offers and gamblers that always redeem offers – we find that reference- dependence sensitivities are similar across all three groups of gamblers.[11] Second, the demographics are very similar between both groups: 34.2% of

[11] Results in the Web Appendix.

redeemers are male, 37.2% of non-redeemers are male and the average age between groups is identical (61). Last, in Table 2 we present four regressions to predict offer redemption and offer value conditional on redemption as for both types of offers. In both regressions we control for gambler fixed effects and the trip number of the gambler. First, we estimate the probability of receiving an offer conditional on the last trip theo and casino win. We find that for both offer types the coefficient casino win is not significant, but the coefficient on theo is significant. This suggests that variation in casino win does not influence offer redemption, but theo does. This aligns with our discussions with the firm that offers are based on theo win, not actual outcomes. In the last two columns we predict the expected value of an offer (conditional on receiving an offer) using linear regression. We see again that the coefficient on players' actual losses is not significant – only theo is correlated with the offer value, regardless of offer type. Once we condition on theo, actual outcomes do not affect either redemption propensity or the value of offers redeemed.

We mention two other caveats on our marketing variables as it relates to our analysis. First, there may be a concern that larger offer values are correlated with larger outcomes (either gains or losses). Because of this, we model the *change* in play volume, not play volume itself. We would expect larger play volumes from larger offers, but the relationship between offer values and the change in play volume between trips is less obvious. Second, we note that for any given gambler the timing of these offers is relatively random: if the timing of the offer does not align with the planned trip or the offer list was prepared when a gambler was in the middle of a trip they gambler might not be included in the offer. This means that some gamblers, even with a qualified last trip theo, simply may not get an offer due to the logistics and timing associated with curating and sending marketing offers.

In sum, based on our conversations with the firm and the evidence in the data, we find that marketing offers are not correlated with actual outcomes to any meaningful degree due to the targeting rules of the firm and the inherent randomness of gamblers' outcomes. We also find no evidence that gamblers who redeem offers are substantially different from gamblers who do not redeem offers, partially because the vast majority of gamblers have a combination of both redemption and non-redemption trips.

3 Model free evidence, reference-dependent model specification, & homogeneous results

To motivate reference-dependent targeting, we provide evidence that player behaviors are sensitive to our three reference points of interest: 1) the current trip outcome (relative to zero), 2) the prior trip outcome, and 3) the expected loss. In Section 3.1 we present model free evidence which suggests gamblers pay attention to these reference points. Section 3.2 formally operationalizes our reference-dependent model. Section 3.3 presents results from a homogeneous model and show the advantage of incorporating reference-dependence over using only customers' RFM information.

Two of the three reference points are directly observed by the gambler (current and prior trip outcomes). The remaining reference point, the expected loss, deserves discussion. While gamblers do not observe theo directly, we use this reference point to serve as a proxy for gamblers who may experience a relatively “lucky” or “unlucky” trip. For instance, if after playing for a relatively long time a gambler nearly breaks even, the gambler would presumably feel that this was a gain relative to the same outcome over a very short period of play. We capture this type of phenomenon by measuring actual outcomes relative to theo. To assess external validity, our conversations with slot machine enthusiasts reveal that some possess a fundamental grasp of theoretical loss, and they anticipate losses over extended periods of play. They apply this cognitive framework to evaluate their gaming experiences. One individual shared, “I go to Atlantic City once a month to play slots, and I budget to lose $5,000 a year. If I did better than that, then I am happy and consider myself ahead, because these trips brought me joy and someone had to pay for all the lights at casinos” Another recounted, “During my recent 10-day cruise, I spent two hours each night at the ship’s casino. I broke even by the end of the trip, which felt like a victory to me.”

At first glance, it may appear that these three reference points are specific to the casino industry but they easily generalize to many other contexts. For example, in the financial services industry (a $3.6 trillion industry), investors experience performance outcomes through their portfolio managers. They can evaluate these outcomes relative to zero, relative to prior outcomes (e.g., prior quarter performance), and relative to expectations (e.g., historical average return of the market or portfolio). Or, in the hospitality industry, a person traveling to a particular hotel can evaluate the service as positive or negative relative to either prior stays with the same hotel (e.g., not as bad as before), or relative to expectations (e.g., based on reviews). The difference with the casino industry, and what makes it ideal for analysis is that the house advantage is known, so we know exactly what *should* have happened based on the machine parameters, and can compare that to what *actually* happened. By analyzing many gamblers, we are able to compare the difference in future behavior based on the actual outcomes experienced relative to prior trips and machine expectations.

3.1 Model free evidence

Our primary dependent variable is the change in theo between trips. We also refer to this as change in play volume. While casino revenue is driven by actual player wins and losses, actual outcomes are essentially random on any given trip. Because of this, casinos tend to focus on increasing play volume, with the understanding that over a large number of gamblers the sum of actual player losses will be similar to the total of theoretical loss. We model the percent change in return trip play volume to accommodate the variation in baseline play across gamblers.[12] Specifically:

[12] As opposed to predicting the expected theo in dollars or the difference in theo between trips.

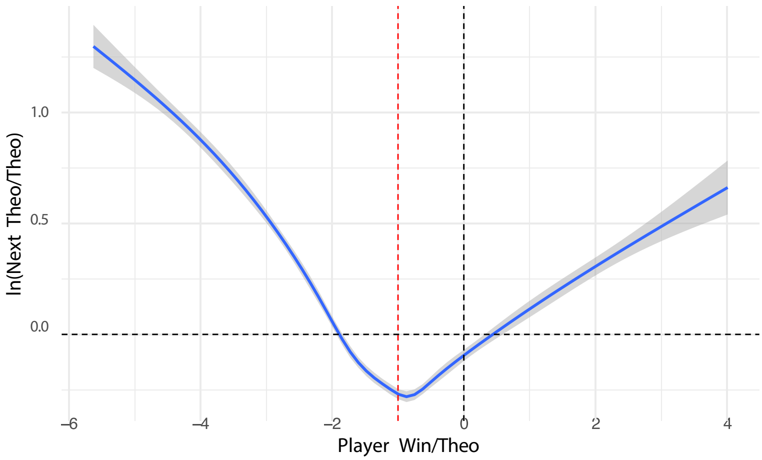

Fig. 2 Change in Return Trip Play Volume vs. Player Win Relative to Theoretical Loss Relationship between actual outcome, theoretical loss, and return trip play volume for trips two through seven in the data (excluding outer 5% of theoretical loss values). For clarity, we hide individual data points and show only the LOESS curve. The horizontal axis shows player win relative to theoretical loss, where positive values are associated with a player win, negative values a player loss. The vertical dashed red line at −1 indicates that the gambler lost exactly what was expected based on their play and house advantage. The vertical axis is the log ratio of next trip play volume relative to current trip play volume. Values greater than zero (above the horizontal dashed line) indicate an increase in play

$$\Delta \text{Theo}_t = \ln\left(\frac{\text{Theoretical Loss}_{t+1}}{\text{Theoretical Loss}_t}\right) \tag{2}$$

where t is the trip number of the current trip. Figure 2 shows a fitted line of changes in play volume against the ratio of player win over theoretical loss. The ratio on the horizontal axis indicates the actual outcome relative to theo (recall theo is always positive – positive ratios indicate a player win). Using a ratio allows us to compare gamblers that play at different levels. To minimize survival bias, in these figures we only consider the trips two through seven, which retains about 44% of all trips (all gamblers have at least seven trips).[13]

The horizontal axis has two key points: 0 (marked by the vertical black dashed line) and −1 (marked by the vertical red dashed line). Any value on the horizontal axis greater than 0 means that the player won, and any value less than 0 means the player ended the trip with a loss. The red dashed line represents the point at which the actual loss equals the expected loss. Even though gamblers never observe their theo level (and it is not provided to them), we see this still serves as a key pivot point of behavior. This may reflect gamblers' expectation of losing at least some money when gambling – breaking even after a long gambling session could be seen as a very positive experience (even if the monetary gain is zero). Points to the right of

[13] In the Web Appendix we show that there is very little difference in estimates when considering the first seven trips as opposed to all trips.

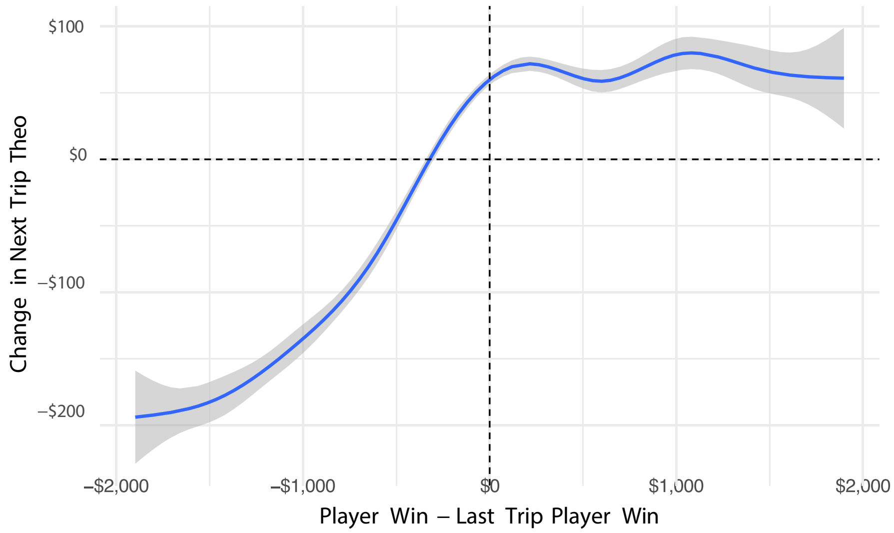

Fig. 3 Change in Return Trip Theoretical Loss vs. Difference in Current and Last Trip Outcome Relationship between outcome relative to the prior trip and subsequent change in play volume on the return trip (excluding outer 5% of theoretical loss values) in trips two through seven. For clarity, we hide individual data points and instead show the LOESS curve. Each data point requires three sequential casino trips: the horizontal axis shows the difference in the gambler outcomes between first two trips, and the vertical axis shows the change in theoretical loss (i.e., play volume) between the second and third trip in the sequence

the red dashed line imply a player was "lucky": they either won money (greater than 0) or lost less than they were supposed to given the house advantage of the machines they played (between -1 and 0). Likewise, points to the left of the red line (less than -1) represent "unlucky" trips: these gamblers lost more than they were supposed to given the house advantage of the game and their playing behavior. The other key point is the horizontal line at 0. Values greater than 0 on the vertical axis imply that the gambler increased their play on the subsequent trip, and values less than 0 represent a decrease in play on the return trip.

Figure 2 presents a few interesting patterns. First, when players pass the break-even point (Player Win/Theo > 0), play begins to increase. That is, winners tend to increase play volume on the return trip.

We see similar behavior for gamblers that are particularly unlucky, starting at a Player Win/Theo value of -2 and lower. Also, we note that the slope is much steeper on the left side of the graph versus the right side, suggesting that (on average) gamblers are more likely to increase play in response to additional losses relative to additional wins.

Second, we see that the expected loss, not actual loss, serves as an important pivot point in play behavior, even though it is not directly observed by the gambler. Next trip play is lowest when a gambler loses what they were expected to lose. One potential explanation for this is that deviations from expected losses (what we term "lucky" and "unlucky" outcomes) could also represent the "excitement" of the gambling session, which is presumably lowest when a gambler losses what is expected from the machine. As the deviations from expected losses increase, return trip play volume increases as well – regardless of the direction of gain or loss.

While Figure 2 allows us to see the interplay between actual and expected losses, in Figure 3 we focus on the remaining reference point: the last trip outcome. The horizontal axis represents the difference between the player win on the current trip relative to the last trip, while the vertical axis shows the change in theoretical loss (i.e., play volume) on the next trip relative to the current trip. Specifically, each point in the graph involves three sequential trips: the relative difference in outcomes between trips one and two, and the relative difference in play volume on trips two and three.

Here again we see that the gain or loss relative to the last trip acts as a key pivot point: gamblers that tend to have a gain relative their last trip tend to increase play volume, but on average it does not vary much with the magnitude of the gain (areas greater than \$0 on the horizontal axis). However, when gamblers begin to have a loss relative to the last trip outcome, change in play volume correlates strongly with the magnitude of the loss. It is important to reiterate that a gain relative to the previous trip does not imply that the player won on the recent trip. For example, a player may lose \$300 on one trip then lose \$200 on the return trip – this is a gain of \$100 relative to the last trip.

The model free evidence suggests that gamblers change behavior substantially in response to gains and losses at different reference points. We see evidence that gamblers pay attention to more than simply the gain or loss from their most recent trip: there is evidence that gamblers also react to gains and losses relative to 1) unobserved expected losses and 2) outcomes from prior trips. The initial findings suggest there might be an opportunity to target on reference-dependent behavior. In the next section we present our approach to estimate these effects.

3.2 Model of reference-dependent sensitivities

To operationalize our three reference points, Table 3 outlines our parameters and associated variables, where c is the current trip gain, l is the last trip gain, and h is the expected loss (which will always be negative). For example, suppose a gambler loses \$400 on the first trip and wins \$800 on the second trip, with an expected loss of \$600 on the second trip. Then, at the end of the second trip we have $c = \$800$, $l = -\$400$, and $h = -\$600$, from which we can construct the relevant gain and loss reference measures.

Table 3 Reference Dependent Parameters and Variables

Reference	Parameter	Variable Description	Variables Used
Current Trip	ω	Player win relative to zero	c
Last Trip	θ	Gain or loss relative to the last trip	c and l
Theo	ρ	Gain or loss relative to expected player loss	c and h

We allow gambler utility to be a function of all three components, with the intent of identifying which heuristics, if any, relate to return trip play volume. For each reference dependence point, we include both an indicator for positive values, and sign dependent magnitudes to capture the individual's sensitivity to gains and losses. More formally, for a given player outcome c, prior trip player outcome l, and theo loss h we set the risk characteristic Γ as:

$$\begin{aligned}\Gamma(\cdot \mid c, l, h, \omega, \theta, \rho) = {} & \omega^I[\mathrm{I}(c > 0)] + \omega^-[\mathrm{I}(c > 0) \cdot c] + \omega^+[\mathrm{I}(c > 0) \cdot c] + \\ & \theta^I[\mathrm{I}(c > l)] + \theta^-[\mathrm{I}(c < l) \cdot (c < l)] + \theta^+[\mathrm{I}(c > l) \cdot (c - l)] + \\ & \rho^I[\mathrm{I}(c > h)] + \rho^-[\mathrm{I}(c < h) \cdot (c - h)] + \rho^+[\mathrm{I}(c > h) \cdot (c - h)] \\ & = [\Phi]'[T]\end{aligned} \tag{3}$$

Thus, for each reference dependence point, we have three coefficients labeled with superscripts: I to indicate if the value is positive, $+$ for the magnitude if positive, and $-$ for the magnitude if negative. The nine coefficients are collected in the column matrix $\Phi = [\omega^I \ \omega^- \ \omega^+ \ \theta^I \ \theta^- \ \theta^+ \ \rho^I \ \rho^- \ \rho^+]$ and the nine associated variables are summarized in the column matrix T.

This formulation is flexible enough to capture a wide range of gain and loss response behaviors relative to each reference point. First, the indicator variables allow for discontinuous jumps at the border of losses and gains. Second, by having separate variables for gains and losses we allow for asymmetric response behavior. For instance, we may find that some gamblers tend to increase play volume as both gains or losses increase. Or we may find that some gamblers respond positively to gains but scale back when encountering losses. Alternatively, gamblers might focus more on the binary outcomes of losses and wins, but are not as sensitive to the magnitude. In short, this formulation in tandem with our detailed casino data will provide greater insight into how individuals' prior outcomes relate to future behavior.

3.3 Homogeneous model specification

To operationalize our risk function Γ proposed in Section 3.2, we incorporate it into a model of gambling utility, which is a function of risk, experience, demographics (age and indicator for male), and factors unobserved by the analyst. The utility for gambler i on trip t is reflected through changes in play volume relative to trip t: Δ Theo$_{it}$. Thus our base specification is as follows:

$$\Delta\text{Theo}_{it} = \alpha + \kappa \text{Trip Number}_{it} + \Psi Z_i + \Gamma(\cdot)_{it} + \varepsilon_{it} \tag{4}$$

where α is the intercept, κ controls for gambling experience at this casino, Ψ the collection of coefficients on our two demographic variables, Γ the risk function previously specified, and ε the variation unobserved to the analyst.

We are aware that play volume on the return trip will be influenced by the outcomes occurring within that trip, but since actual outcomes from a single trip are essentially normally distributed across players, we instead focus on conditioning on gains and losses from the prior trip. Of managerial interest is the drivers of unknown return trip play volume using marketing actions and information of gamblers' last

Table 4 Homogeneous results

	Dependent variable:		
	Δ Theo		
	Gain/Loss Heuristics	RFM	Both
Intercept	0.233*** (0.022)	0.234*** (0.023)	0.092*** (0.022)
Player Win + Indicator	0.278*** (0.014)		0.286*** (0.014)
Player Win +	0.0003*** (0.00002)		0.001*** (0.00002)
Player Win -	0.001*** (0.00001)		0.001*** (0.00001)
Rel. Last Trip + Indicator	0.379*** (0.010)		0.370*** (0.009)
Rel. Last Trip +	0.0002*** (0.00001)		0.0001*** (0.00001)
Rel. Last Trip -	– 0.0001*** (0.00001)		– 0.0001*** (0.00001)
Rel. Theoretical Loss + Indicator	– 0.395*** (0.010)		– 0.402*** (0.010)
Rel. Theoretical Loss +	– 0.001*** (0.00001)		– 0.001*** (0.00001)
Rel. Theoretical Loss -	– 0.0002*** (0.00002)		– 0.0004*** (0.00002)
Trip Number	– 0.001*** (0.0003)	– 0.002*** (0.0004)	0.001 (0.0004)
Age	0.002*** (0.0003)	0.001** (0.0004)	0.002*** (0.0003)
Male	– 0.025*** (0.007)	– 0.017** (0.008)	– 0.017** (0.007)
Days Since Last Trip		0.0001*** (0.00001)	0.0001*** (0.00001)
Days Between Trips		– 0.0001*** (0.00002)	0.0001*** (0.00002)
Avg. Theo		– 0.0005*** (0.00001)	0.001*** (0.00001)
Observations	142,212	142,212	142,212
Adjusted R^2	0.135	0.020	0.148
RMSE In-Sample	1.322	1.417	1.322
RMSE (5 Fold CV)	1.322	1.418	1.322

$^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

trip outcomes. As most casino marketing managers are data- and time-constrained and thus unable to market to customers in a real-time fashion within the current trip, we believe our approach of linking past trip outcomes to return trip play volume is both externally valid and managerially actionable.

The first column in Table 4 shows the results from the estimation of the base model with homogeneous parameters.[14] All nine risk heuristics are significant. Recall that the outcome variable is the log of the ratio of the return trip play volume relative to the prior trip play volume – positive expected values indicate an increase in play volume and negative expected values indicate a decrease in play volume. The intercept is positive but the coefficient on trip number is negative, suggesting that in general play volume increases on return trips but at a decreasing rate. A positive coefficient on age indicates that the increases from older gamblers is larger. Similarly, the negative coefficient on male suggests that females tend to show greater increases in play, relative to males.

The coefficient on the player win indicator is positive – gamblers tend to increase their play if the trip resulted in a gain relative to zero, regardless of magnitude (Player Win + Indicator). We also see a positive coefficient on the *magnitude* of wins (Player Win +), suggesting that future trip play volume increases in tandem with current trip wins. The coefficient on player losses (Player Win -) is also positive, implying that gamblers reduce their play when faced with a loss. Interestingly, the sensitivity on the magnitudes of losses is about three times the magnitude on the sensitivity on gains, which aligns with previous research on loss aversion.

The second reference point considered is the last trip outcome (Rel. Last Trip). The indicator variable is about 1.5 times the magnitude of the coefficient on the player win indicator, suggesting that a gain relative to the last trip is associated with a greater increase in play versus a gain relative to zero on the current trip. The coefficient on the magnitudes of these gains is also positive, but for losses it is negative. This suggests that gamblers tend to increase play as the magnitude of the loss increases relative to the last trip outcome.

Finally, we consider the third reference point: gains and losses relative to the theoretical loss (Rel. Theoretical Loss). Again, the player does not directly observe their theoretical loss, but we posit that a gambler will understand to some extent whether they were particularly lucky (or unlucky) in a given trip given their actual outcome, time played, and average bet. Therefore, this serves as our indicator of the player luck. Importantly, a player can be lucky but still have a negative outcome (for example if a gambler loses \$100 but was expected to lose \$1,000 based on their theo). Being "lucky" but facing a loss is akin to a player enjoying the gambling experience for an extended time at very low cost – analogous to a very low house advantage on the machine. The "relative to theo" coefficients suggest that when a player is lucky (i.e., a gain relative to the theoretical loss), they tend to have a significant pullback in play. In regards to the magnitude, the reduction in play is greater the "luckier" the gambler is (that is, when the outcomes are more positive than what was expected).

[14] We estimate the model on all trips, but find similar results when using trips two through six alone to account for potential selection bias. These results are available in the Web Appendix. In addition, to rule out time dependence, we also interacted each heuristic with trip number to capture any potential time varying nature of the coefficients but this does not explain any additional variation in return trip play volume. Results are available by request.

Interestingly, we see that unlucky gamblers (Rel. Theoretical Loss -) tend to increase play, though to a lesser extent than lucky players.

In short, we see evidence for numerous asymmetric effects which vary based on the heuristic considered. To ensure that our added heuristics are simply not overfitting the data, we include the root mean-squared- error both in-sample and with 5-fold cross validation.[15] By the nature of how these variables are derived we are not concerned with multicollinearity: since casino outcomes are random, we do not expect strong correlations between theo and actual outcomes mostly because an individual trip is too short for any law of large numbers to manifest itself in a meaningful way. By the same logic, these correlations would be weak when compared to prior trip outcomes. Similarly, the three variables within each reference point again would not be correlated: an indicator by construction will not be correlated with continuous levels. Finally, due to the random nature of casino outcomes it unlikely that positive and negative outcomes will be correlated in any way.[16]

3.3.1 Comparison to RFM performance

The last two columns of Table 4 consider the performance relative to a standard recency, frequency, monetary (RFM) model. The results show that the gain and loss heuristics substantially improve the fit. First, we see that the percentage of variation explained increases by more than a factor of three when the nine heuristics are included. This suggests that the heuristics are providing much more insight into variation in next trip play volume relative to traditional RFM metrics currently used in the industry. Second, we see that RMSE, both in-sample and cross-validated, decrease with the additional heuristics, attesting to their contribution of predictive power. These results indicate that traditional RFM metrics are relatively insufficient in explaining return trip play behavior (at least in this empirical context) and that casinos might instead improve their understanding of the drivers of play volume by incorporating the gains and losses relative to the reference points we propose.

4 Marketing interactions & heterogeneity

The model-free and homogeneous results provide evidence that gambler behavior is sensitive to gains and losses relative to the three reference points proposed. The findings motivate a potential opportunity for firms to target on reference-dependent behavior. What remains to be seen is whether marketing actions interact with these sensitivities. Firms interested in targeting customers who meet specific gain/loss

15 We also tried tenfold cross-validation with similar results.

16 Still, we performed variance inflation factor check on our models and the VIF never exceeds 5, with most VIFs uniformly distributed between 1 and 3.

profiles need to be aware that the marketing actions themselves may influence both play volume and sensitivities to gains and losses.[17]

To account for the effects of marketing, we expand our model by interacting our two marketing variables (room comps and promo credits) with the nine gain/loss heuristics. In addition, we allow for full heterogeneity across individuals through a hierarchical Bayesian model. The full specification with marketing interactions and individual heterogeneity is as follows:

$$\begin{aligned} \Delta\text{Theo}_{it} &= \alpha_i + \mathcal{K}_i \text{Trip Number}_{it} + \left[\Phi_i \ \Upsilon_i \ \Xi_i\right]\left[[T]\left[1 p_{it+1} \ r_{it+1}\right]\right] + \beta_i^p p_{it+1} + \beta_i^r r_{it+1} + \varepsilon_{it} \\ \eta_i &\sim N\left(Z\psi_{[i,]}, V_\eta\right) \end{aligned} \tag{5}$$

Where 1 is the unit value, p_{it+1} is the value of the return trip promo credit offer for gambler i after trip t, and r_{it+1} the value of the return trip room comp offer. T is the vector of nine gain/loss variables, as specified earlier. The matrix $[\Phi_i \ \Upsilon_i \ \Xi_i]$ is the collection of individual level coefficients associated with the main effects, interactions with promo credits, and interactions with room comps, respectively. Thus in total there are 20 additional parameters: 18 for each of the 9 interactions with promo credits and room comps, and 2 for the main effects of each. That is, we take all elements of the original risk function and interact with return trip promo credits and return trip room comps. All individual-level coefficients are collected in η_i.

In the second level of our specification, we allow for a normal distribution of heterogeneity.[18] Z contains an intercept and the mean-centered demographic variables (age and indicator for male).[19] $Z\psi_{[i,]}$ refers to the ith row of the product of the demographic variables Z and ψ. We use the offer value on the return trip ($t + 1$) for two reasons. First, our emphasis is on changes in return trip behavior, so we need to account for variables that might influence return trip play. The second reason is for operational reasons: this is what marketers have control over. Once the gambler is back on property, any offer used for that trip is already a sunk cost.

The three columns under "Average Gambler" in Table 5 shows the median posterior coefficients of the intercept in ψ across gamblers with the lower and upper bounds of a 95th highest posterior density (HPD). As expected, the posterior median across gamblers is similar to homogeneous results presented earlier: a positive coefficient on player win and last trip reference, but a negative coefficient on gains relative to theoretical win.

The last three columns of Table 5, under "Across Gamblers", reflect the variation across gamblers in their individual posteriors. As is standard practice with our specified hierarchical linear model structure, the demographics will moderate the intercept in ψ to generate individual-level estimates η. In the table we present the standard deviation in the posterior medians of η across gamblers, along with the 25th and 75th percentiles.

17 We are unable to disentangle whether marketing offers moderate the gain and loss sensitivities, or whether the state of the gambler conditional on their gains and losses moderates sensitivities to marketing. For purposes of this analysis this distinction is not necessary.

18 We tried up to a five component mixture of normals on the heterogeneity distribution, with similar results.

19 Please see the Web Appendix for the full hierarchical structure and prior specifications.

Table 5 Hierarchical Bayesian Model Results with Marketing and Reference-Dependence Interactions

	Average Gambler (ψ)			Across Gamblers (η)		
	Median	Lower	Upper	Std. Dev	25th	75th
Intercept	0.7484***	0.7056	0.7981	0.1054	0.6941	0.8084
Trip Number	0.0345***	0.0295	0.0395	0.0273	0.0271	0.0447
Player Win + Ind. (PWI)	0.2557***	0.1822	0.3214	0.0823	0.2165	0.2845
Player Win + (PW +)	0.0048***	0.0020	0.0070	0.0121	0.0011	0.0077
Player Win-(PW-)	0.0056***	0.0045	0.0066	0.0066	0.0019	0.0079
Rel. Last Trip + Ind. (LTI)	0.2741***	0.2374	0.3176	0.0773	0.2558	0.2980
Rel. Last Trip + (LT +)	0.0002	−0.0007	0.0013	0.0050	−0.0015	0.0017
Rel. Last Trip-(LT-)	−0.0010	−0.0021	0.0001	0.0057	−0.0028	0.0010
Rel. Theo Loss + Ind. (TLI)	−0.3072***	−0.3682	−0.2534	0.0826	−0.3501	−0.2656
Rel. Theo Loss + (TL +)	−0.0050***	−0.0063	−0.0036	0.0087	−0.0076	−0.0012
Rel. Theo Loss-(TL-)	−0.0044***	−0.0058	−0.0033	0.0101	−0.0082	0.0002
Promo Credits (1 k)	1.2267***	0.1025	1.5191	0.0634	1.0320	1.1040
Room Comp (1 k)	1.7734***	1.3990	1.9580	0.2803	1.5584	1.9193
PWI x Promo Credits (1 k)	1.1353	−0.0649	1.6078	0.8677	0.0532	1.6782
PW + x Promo Credits (1 k)	−0.0038	−0.0198	0.0130	0.0181	−0.0100	0.0023
PW- x Promo Credits (1 k)	−0.0046***	−0.0080	−0.0012	0.0122	−0.0087	0.0003
LTI x Promo Credits (1 k)	0.1772	−0.3130	0.5147	0.2368	−0.1547	0.3045
LT + x Promo Credits (1 k)	0.0036	−0.0003	0.0079	0.0143	−0.0005	0.0069
LT- x Promo Credits (1 k)	−0.0007	−0.0049	0.0033	0.0138	−0.0048	0.0038
TLI x Promo Credits (1 k)	0.0800	−0.1801	0.5305	0.2264	−0.0196	0.2779
TL + x Promo Credits (1 k)	0.0039	−0.0031	0.0091	0.0137	−0.0004	0.0076
TL- x Promo Credits (1 k)	−0.0011	−0.0080	0.0055	0.0164	−0.0060	0.0042
PWI x Room Comp (1 k)	0.1923	−0.0043	0.5713	0.2600	0.0742	0.4321
PW + x Room Comp (1 k)	−0.0033	−0.0130	0.0060	0.0134	−0.0081	0.0013
PW- x Room Comp (1 k)	−0.0020	−0.0047	0.0004	0.0086	−0.0049	0.0016
LTI x Room Comp (1 k)	−0.1274	−0.2998	0.1256	0.0906	−0.1619	−0.0523
LT + x Room Comp (1 k)	−0.0003	−0.0032	0.0023	0.0093	−0.0033	0.0025
LT- x Room Comp (1 k)	0.0005	−0.0020	0.0031	0.0093	−0.0027	0.0038
TLI x Room Comp (1 k)	0.4947***	0.2910	0.7958	0.3387	0.2958	0.7220
TL + x Room Comp (1 k)	0.0022	−0.0015	0.0063	0.0108	−0.0014	0.0056
TL- x Room Comp (1 k)	0.0001	−0.0039	0.0039	0.0127	−0.0037	0.0044

* HPD at 90% excludes 0; ** HPD at 95% excludes 0; *** HPD at 99% excludes 0

There tends to be agreement on the indicator coefficients (I), but much more dispersion in the coefficients associated with magnitudes (+ or −). For example, at the median gains relative to the last trip are associated with increases in return trip play volume, but the standard deviation suggests that for many gamblers greater gains will instead result in decreases in play volume on the return trip. While many of the marketing interactions are not significant for the average gambler, in our targeting exercise we use both these average effects and individual estimates.

5 Projected targeting improvements

In this section we evaluate the effectiveness of reference-dependent targeting by comparing profits across five casino targeting policies. The first policy, "Theo Only", reflects the status quo of most casinos where targeting is based solely on past theoretical loss and does not account for actual outcomes or reference- dependent sensitivities. The other four targeting policies rely on predictions from various hierarchical linear model (HLM) specifications. Let $\pi(x_{it}; m)$ indicate whether gambler i at trip t is targeted under model specification m:

$$\pi\left(x_{it};m\right) = \mathrm{I}\left(\mathbb{E}[y_{it+1}(1) - y_{it+1}(0)|X = x_{it}^{m}] > c\right) \tag{6}$$

The gambler is targeted if the difference in expected theo between using an offer on the return trip, $y_{it+1}(1)$, and not using it, $y_{it+1}(0)$, exceeds the cost of the offer c. Since not all targeting policies use the same variables for predictions, x^m indicates the relevant subset of of data associated with model m.

Table 6 reports expected difference in gambler profits across targeting policies for a standard room comp and promo credit offer. Predicted changes in profits are based on a separate, flexible hierarchical linear model that incorporates all variables and interactions used across policies; we denote this set of variables as x^*_{it}. We calculate the expected change in gambler-trip profit as follows:

$$\Pi_{(m)} = \frac{1}{N_{\pi}(m)} \sum_{i.t} \pi\left(x_{it};m\right) \times \left(\mathbb{E}\left[y_{it+1}(1) - y_{it+1}(0)\middle|X = x_{it}^{*}\right] - c\right) \tag{7}$$

where $N_{\pi}(m) = \sum_{i,\,t} \pi_{it}(m)$ is the total number of gambler-trips targeted under model m. We first consider a standard room comp offer valued at \$300. The assumed cost of this offer is \$30, which reflects the labor cost of preparing a room for one night – we ignore the opportunity cost of selling the room elsewhere. To evaluate the "Theo Only" policy, we need a reasonable threshold for when a gambler-trip qualifies for an offer. From the data we find that the room comp reinvestment rate is about 5%: every \$100 in theoretical loss earns \$5 towards a room comp. Therefore, a gambler-trip which generates \$30/5% = \$600 in theoretical loss would be eligible for a \$300 room comp offer. Under this targeting policy, the casino earns an average of \$69 in profit.

In our benchmark model ("Individual (RFM Only)") we include trip number, last trip play volume, days since the last trip (R), trip frequency (F), and average trip play volume (M) but no reference-dependent variables. As with our reference-dependent specification, we interact the RFM variables with marketing offers and estimate individual level coefficients. The RFM model yields \$283 in profit.

The remaining three policies use our specified reference-dependence structure in varying levels of com- plexity: "Average gamblers" uses only the intercept values ψ of each variable, reflecting a situation where we have a new gambler and not enough data or demographics to estimate precise individual-level posteriors. The second variation adds demographics. The third variation uses the full posteriors with

Table 6 Expected Average Difference in Profit Per Gambler

	$300 Room Comp		$100 Promo Credit	
	% Targeted	Avg. (SE) Profit	% Targeted	Avg. (SE) Profit
Theo Only	29%	68.98 (2.05)	30%	30.19 (1.30)
HLM				
Individual (RFM Only)	68%	283.13 (3.18)	36%	260.70 (4.37)
Average Gambler	77%	228.49 (2.49)	24%	50.33 (1.57)
Demographics	77%	228.95 (2.49)	23%	50.28 (1.57)
Individual (Reference-Dependence Only)	70%	339.26 (3.90)	66%	397.87 (5.12)

Expected average profit difference per targeted gambler-trip. Profit differences are based on the conditional average treatment effect of using an offer among those targeted, with the standard error of the mean shown in parentheses

all individual-level coefficients, as specified in equation 5 (“Individual (Reference-Dependence Only)”).

The HLM using the average gambler coefficients performs about the same as the HLM with demographics ($228 and $229, respectively), indicating, similarly to Rossi et al. (1996) and Smith et al. (2023), that demographics do not provide substantially valuable information for targeting purposes. Still, we only have two demographic variables (age and an indicator for males) – richer demographic information may yield better gains. Our reference-dependent HLM with individual-level heterogeneity outperforms all other targeting policies, yielding $339 in profit, and increase of 19.8% over the benchmark RFM model.

We now consider the promo credit offer. In our data, the average reinvestment rate for promo credits is about 17.5%, implying that for the “Theo Only” policy, as soon as a gambler generates around $571 in theoretical loss they qualify for a $100 promotional credit. Promotional credits are usually treated as a “hard cost” so we use the full $100 as the cost of this offer.[20] As with the room comp offer, the theo-based targeting policy underperforms since it does not consider projected changes in play volume, only historical expected losses, and yielding only $30 in profit.

The individual-level RFM model resulted in substantially higher profits ($261) relative to the theo- based policy. In our reference-dependent specification, targeting based on the average gambler performs nearly identically to the targeting policy that incorporates demographics, both yielding only about $50 in profit. The individual-level reference-dependent specification yielded $398 in profit, outperforming the RFM benchmark by 52.6%

[20] Since promo credits can be cycled back into the machine, the actual cost to the casino can vary.

6 Discussion and conclusion

Using transaction data from the casino industry and incorporating individual reference-dependent sensi- tivities into CRM practices, we offer a novel approach to enhancing targeted marketing efficiency. The contribution of our research lies in the exploration of heterogeneous reference-dependent sensitivities and advancing our understanding of their effects using real-world data. Managerially, our study offers valu- able insights for the casino industry. By incorporating individual-level reference-dependent sensitivities into casino targeting strategies we estimate that targeting profitability increases by at least 19.8% relative to an RFM model without reference-dependent sensitivities, depending on the offer type. These findings underscore the potential benefits of incorporating sensitivities to gains and losses at different reference points into customer relationship management strategies.

Our findings have implications for the casino industry and many other service domains where under- standing customers' experiences relative to reference points can be leveraged for engagement strategies. For example, in the financial services industry, clients experience fluctuations in the quarterly performance of their portfolios. These quarterly fluctuations may be greater or less than 1) zero, 2) the prior quarter, or 3) their expectations (e.g., historical market averages). The firm managing investor relationships needs to determine how best to allocate their CRM resources given the gains and losses facing each client relative to each reference point. Other industries that share similar properties include travels, hotels, restaurants, delivery services (Gui & Drerup, 2022) and ridesharing services.

Our modeling focus of engagement conditional on a visit can be extended to a wide range of non- contractual service industries where companies have limited control over the frequency of customer visits but substantial influence over customer engagement during those visits. For instance, firms operating in the hospitality industry, hotels, resorts, cruiselines and amusement parks have limited control over when people choose to travel due to various personal, professional, and economic constraints. However, targeted marketing strategies informed by past customer experiences can allow these businesses to significantly enhance guest engagement and spending for when they do return.

In addition, our insights can be applied to contractual business settings where the timing of customer interactions may be fixed yet the scope for enhancing user engagement and optimizing product usage remains relevant. In industries like subscription-based services, including monthly delivery boxes, SaaS, and digital education platforms, companies can refine their CRM strategies by focusing on usage patterns and customer experiences rather than the frequency of visits or purchases. For example, companies offering subscription boxes can utilize data on past customer experiences to affect future purchase contents and amounts. Similarly, SaaS providers benefit from monitoring how users interact with their software. Educational platforms that operate on subscription models can benefit from understanding and reacting to user engagement levels. By recommending new courses based on previously completed ones or adjusting learning pathways according to user experience, these platforms can improve user engagement and educational outcomes.

Despite its contributions, our study is not without limitations and provides avenues for future research. First, we do not consider especially sophisticated machine learning models. This is an active choice of ours given the scope of the analysis in testing the presence and relative strength of various reference-dependent heuristics, as well as feedback from our industry contacts on current practices. Clearly, a more sophisticated method incorporating reference-dependent heuristics should perform even better, so our analysis serves as a conservative benchmark for the potential performance lift in CRM. However, given the limited num- ber of variables analyzed our relatively simple theory-based model specification is likely sufficient. Future researchers with access to more varied customer-level and contextual data can leverage machine learning techniques to improve predictions in casino CRM, potentially uncovering more complex interactions and relationships. Another limitation of our study is the lack of reliable information on customer demographics in the current dataset. Consequently, we are unable to find demographics-dependent loss aversion found in experimental research, but this may also be due to the empirical context. Future work can supplement transaction data with more detailed customer demographic information and even survey data to test de- mographic and experience moderators of reference-dependence, allowing for the development of even more personalized marketing approaches.

Finally, it is important to emphasize the potential ethical implications of leveraging individual reference- dependence in marketing practices, particularly in industries where customers are exposed to significant financial risks such as casinos.[21] While our research highlights the potential benefits of incorporating indi- vidual reference-dependent sensitivities into CRM strategies, it is crucial to recognize the potential for misuse or exploitation. In casino marketing, and more broadly in marketing in general, future research could explore the ethical boundaries of utilizing psychological insights in CRM and targeting, ensuring that strategies are not unethically taking advantage of consumers' cognitive biases and leading to negative outcomes such as addiction.

In light of these limitations, additional research can explore the interplay between reference-dependence and other factors that may influence behaviors, such as gambling motivations, impulsivity, or other psycho- logical states. Another promising direction for future research is the examination of the temporal dynamics of reference-dependence. Our study focuses on individual reference-dependence at a specific point in time, but it is possible that customers' sensitivities evolves as they gain more experience with gambling or as their financial circumstances change. Finally, we did not consider the mechanisms as to *why* certain gamblers used some reference points and not others – whether this is due to systematic differences can provide greater theoretical insight to reference-dependence more broadly but is outside the scope of our analysis. Related, it was not clear whether the marketing offers influenced the sensitivities to gains and losses, or whether the state of the gambler after a gain or loss influenced their sensitivity to marketing. Future work may be able to distinguish between these effects.

[21] It is important to note that Narayanan and Manchanda (2012) found that only 7% of all gamblers exhibited evidence of addiction.

Funding Open access funding provided by SCELC, Statewide California Electronic Library Consortium.

Data availability This data will not be available in a public repository due to a non-disclosure agreement between the corresponding author and the firm providing the data. We are open to providing summary or synthetic data and code to ensure transparency in the analysis.

Declarations

Conflicts of interest There are no known potential conflicts of interest.

References

American Gaming Association (2022). State of the states 2022: The aga survey of the casino industry. https://www.americangaming.org/resources/state-of-the-states-2022/. Accessed 10 Jan 2024.

Bleier, A., Goldfarb, A., & Tucker, C. (2020). Consumer privacy and the future of data-based innovation and marketing. *International Journal of Research in Marketing, 37*(3), 466–480.

Brown, A. L., Imai, T., Vieider, F. M., & Camerer, C. F. (2024). Meta-analysis of empirical estimates of loss aversion. *Journal of Economic Literature, 62*(2), 485–516.

Deng, Y., & Mela, C. F. (2018). Tv viewing and advertising targeting. *Journal of Marketing Research, 55*(1), 99–118.

Dew, R., Ascarza, E., Netzer, O., & Sicherman, N. (2023). Detecting routines: Applications to ridesharing customer relationship management. *Journal of Marketing Research*, 00222437231189185.

Dong, X., Manchanda, P., & Chintagunta, P. K. (2009). Quantifying the benefits of individual-level targeting in the presence of firm strategic behavior. *Journal of Marketing Research, 46*(2), 207–221.

Dubé, J.-P., & Misra, S. (2023). Personalized pricing and consumer welfare. *Journal of Political Economy, 131*(1), 131–189.

Gui, G. and Drerup, T. H. (2022). Designing promises with reference-dependent customers: The case of online grocery delivery time. Available at SSRN 4298782.

Hendel, I., & Nevo, A. (2013). Intertemporal price discrimination in storable goods markets. *American Economic Review, 103*(7), 2722–2751.

Kőszegi, B., & Rabin, M. (2006). A model of reference-dependent preferences. *The Quarterly Journal of Economics, 121*(4), 1133–1165.

Kőszegi, B., & Rabin, M. (2007). Reference-dependent risk attitudes. *American Economic Review, 97*(4), 1047–1073.

Kumar, V., Sriram, S., Luo, A., & Chintagunta, P. K. (2011). Assessing the effect of marketing investments in a business marketing context. *Marketing Science, 30*(5), 924–940.

Lemon, K. N., & Verhoef, P. C. (2016). Understanding customer experience throughout the customer journey. *Journal of Marketing, 80*(6), 69–96.

Mrkva, K., Johnson, E. J., Gächter, S., & Herrmann, A. (2020). Moderating loss aversion: Loss aversion has moderators, but reports of its death are greatly exaggerated. *Journal of Consumer Psychology, 30*(3), 407–428.

Nair, H. S., Misra, S., Hornbuckle, W. J., IV., Mishra, R., & Acharya, A. (2017). Big data and marketing analytics in gaming: Combining empirical models and field experimentation. *Marketing Science, 36*(5), 699–725.

Narayanan, S., & Manchanda, P. (2012). An empirical analysis of individual level casino gambling behavior. *Quantitative Marketing and Economics, 10*, 27–62.

Padilla, N., Ascarza, E., and Netzer, O. (2023). The customer journey as a source of information. Available at SSRN 4612478.

PlayUSA (2022). United states casino statistics. https://www.playusa.com/united-states-casino-statistics/. Accessed 1 Feb 2024.

Rossi, P., McCulloch, R., & Allenby, G. (1996). The value of purchase history data in target marketing. *Marketing Science, 15*(4), 301–394.

Simester, D. I., Sun, P., & Tsitsiklis, J. N. (2006). Dynamic catalog mailing policies. *Management Science, 52*(5), 683–696.

Smith, A. N., Seiler, S., & Aggarwal, I. (2023). Optimal price targeting. *Marketing Science, 42*(3), 476–499.

Taylor, W., & Hollenbeck, B. (2021). Leveraging loyalty programs using competitor based targeting. *Quantitative Marketing and Economics*, 1–39.

Tuchman, A. E., Nair, H. S., & Gardete, P. M. (2018). Television ad-skipping, consumption complementarities and the consumer demand for advertising. *Quantitative Marketing and Economics, 16*(2), 111–174.

Zantedeschi, D., Feit, E. M., & Bradlow, E. T. (2017). Measuring multichannel advertising response. *Management Science, 63*(8), 2706–2728.

Zhang, J. Z., Netzer, O., & Ansari, A. (2014). Dynamic targeted pricing in b2b relationships. *Marketing Science, 33*(3), 317–337.