Incorporating Experience Quality Data into CRM Models: The Impact of Gambler Outcomes on Casino Return Times

Wayne Taylor

September 27, 2016

Abstract

I consider the direct marketing targeting problem in situations where 1) the customer's experience quality level varies from occasion to occasion, 2) the firm has measures of these quality levels, and 3) the firm can customize marketing according to these measures and the customer's behaviors. A primary contribution of this paper is a framework and methodology that allows the manager to assess the marketing response of a forward-looking customer with any specific experience and behavior history, which in turn can be used to decide which customers to target for marketing. This research introduces a novel, tractable way to estimate and introduce flexible heterogeneity distributions into Bayesian learning models with forward looking agents. The model is estimated using data from the casino industry, an industry which generates more than $60 billion in U.S. revenues but has surprisingly little academic, econometric research. The counterfactuals offer findings on gambler learning and direct marketing responsiveness and show that casino profitability increases substantially when marketing incorporates gamblers' beliefs and past outcome sequences into the targeting decision.

Keywords: *dynamic discrete choice, learning models, Bayesian estimation, CRM, targeted marketing*

This paper focuses on the direct marketing problem of whom to target and with what offers in situations where the customer learns about the firm through multiple interactions. Specifically, these situations have the following four characteristics: (a) The customer's experience quality level is random in that it varies from occasion to occasion according to a certain distribution. (b) The customer uses the experiences to learn and form beliefs about the nature of the distribution. These beliefs affect the customer's expected utility from future interactions with the firm, which in turn affect future decisions on whether to interact with the firm. (c) The firm has access to measures of the customer's experience quality levels. (d) The firm can use information on these measures to make different offers to different customers. These four characteristics hold in a large number of industries, like the airline, financial services, restaurant or casino industries. To take an example from the financial services industry: A private banking advisor's performance is often random. A customer may use the performance stream to assess the account's likely long-run return and volatility, which may then influence his/her decision on whether to continue doing business with the firm. The firm observes the performance stream and can make different offers to different customers, potentially customizing according to the performance stream experienced by a customer and according to the customer's behaviors in response to those performances. In airlines, flight delays are random and firms typically record the delays experienced by each passenger. Whenever a customer uses Uber, the driver-rider marketplace firm, the quality of the driver they receive is random. Passengers rate the quality of their experience which Uber can then use for targeting purposes.

A primary contribution of this paper is a framework and methodology that allows the manager to assess the marketing response for any customer with any specific experience and behavior history, which in turn can be used to decide which customers to target for marketing. An important insight identified from this methodology is that the optimal targeting decision may depend on a customer's expectations of quality in addition to their marketing responsiveness. For instance, a customer with early experiences that are of atypically low levels will likely form a belief distribution with a low expectation of experience quality from future visits and, on the basis of this belief, may reduce or altogether stop interactions with the firm. Direct marketing offers targeted at such a customer can increase the customer's expected utility for a future visit, incentivize the customer to interact further with the client, improve the belief distribution based on the new experience, increase the likelihood of future visits, and increase the future profits of the firm from that customer. This suggests that if the marketing offers are costly and the firm can make the offers to only a limited set of customers, then it should target those customers whose future profits will be most increased by the offer, possibly those customers whose belief distributions can be improved the most, which may be those customers who have had atypically low levels of experience quality. However, this intuition should not be taken to imply that the firm should simply direct offers to the customers with the lowest experience qualities.

One has to balance the cost of an offer against the benefits, which will depend on marketing responsiveness and the extent to which the offer will influence the customer's future behaviors and the firm's profit from those behaviors. In this paper, I present a model to do exactly that and provide evidence that incorporating a customer's beliefs of experience quality into the targeting decision can significantly increase a firm's profit.

To address the direct marketing decision problem in this important class of situations where the manager can react to a customer's observed experience and behavior history, I employ a Bayesian, dynamic learning framework. A customer starts with a prior belief distribution on the mean value of the experience quality level and updates this belief distribution according to new experiences accumulated with each interaction with the firm. For a rational consumer, the decision to engage in repeat interactions with the firm, thereby forming more accurate beliefs, will depend on the utility from the interactions, the value of the increased belief accuracy, and the consequent utilities from potential future visits. The model presented in this paper allows one to determine the extent to which direct marketing influences these utilities, the customer's consequent interactions with the firm and firms' consequent profits.

Because prior beliefs, marketing responsiveness and utility function parameters can vary from customer to customer, it is important that the model and estimation methodology allow for across-customer heterogeneity. Incorporating learning into a dynamic discrete choice model with forward looking customers is difficult because the optimal choice is the solution to a complex Bellman equation, with a correspondingly difficult likelihood function. Including unobserved heterogeneity makes the problem even more difficult. An important contribution of this paper is that it develops a tractable solution to this class of problems by combining a forward simulation algorithm with Markov Chain Monte Carlo ideas.

I illustrate this paper's framework and methodology in the context of the casino industry. I focus on the gaming industry for a few reasons. First, casino gambling is a substantial industry in the United States. Gaming revenues are now at an historic peak at $67 billion, with nearly 1,000 casinos operating in 39 states (Oxford Economics, 2014). Oxford Economics estimates the U.S. gaming industry contributes nearly $240 billion to the national economy (2014).[1] The pervasiveness of the industry is as significant as its size; nearly one third of Americans have gambled at a casino within the past twelve months. Therefore, providing insight into this one industry is meaningful. Second, there is limited research on the impacts of casino marketing. This is due primarily to the difficulty of obtaining sensitive casino data rather than the lack of importance of marketing in the gaming industry. Third, it is behaviorally interesting because these are real gamblers responding to uncertainty, as opposed to lab participants. Fourth, an advantage of studying casinos is that gambling outcomes are exogenous. Even though casinos can control the overall house advantage and its

[1]This includes both tribal and commercial casinos. In 2012 71.6 million gamblers spent $37.34 billion in gaming revenues at commercial casinos. This is more than they did on movies, craft beer, and outdoor equipment combined. (American Gaming Association, 2013).

distribution, the variation in trip-to-trip outcomes across individuals is very high. This means that each gambler faces a very distinct and randomly generated outcome sequence, which aides model identification. Finally, many casinos base marketing offer values on gamblers' past *expected* losses, and not their actual outcomes. I discuss this later in more detail but the implication is that much of the offer endogeneity is removed.

I use data on real gamblers to understand how their outcomes influence the time until their return trip. Specifically, I answer the following: how do past outcomes influence gamblers' beliefs on the house advantage and how can marketers use this information in their one-to-one targeting decisions? To estimate these impacts I specify a dynamic learning model that incorporates the belief uncertainty into the utility specification. In the traditional random utility framework, consumers know the attributes of their choices perfectly. Learning models extend the traditional framework by recognizing that consumers may have incomplete information and thus make choices based on perceived rather than actual attributes (Ching et al., 2013). In this model, gamblers learn about the casino's house advantage by gambling at slot machines. Gamblers use their beliefs to form future cost expectations, which influences their decision to return to the casino. The uncertainty in these beliefs can also influence the decision to return. By fully modeling the learning process (rather than simply conditioning on the last outcome) the model permits gamblers to use their entire trip history when forming future cost expectations. The reduced form evidence supports a full structural model of the learning process.

Marketing has shown considerable interest in decision making under uncertainty for decades, roughly starting with Tversky and Kahneman (1974). Meyer (1981) found that temporal variability increases the cost of information gathering, which suggests that variability comes with a premium. More recent work studies uncertainty in customer satisfaction (Bolton, 1998), service quality (Rust et al., 1999; Boulding et al., 1993), and product attributes (Erdem and Keane, 1996). While related, the primary difference is that the focus of this research is on how a customer's belief (and the uncertainty of this belief) can influence a firm's targeted marketing actions. The intent is similar to that of Narayanan and Manchanda (2009), who also study learning behavior and its impact on targeting, however a substantial econometric contribution of this research is that I allow gamblers to be forward looking, which they assume away. Gamblers can make tradeoffs between today's knowns and tomorrow's opportunity to learn more.

The dataset used in model estimation comes from a large casino in the United States. I observe the complete trip histories and marketing activity for a random sample of gamblers. The empirical strategy takes two parts. First, I show descriptive and reduced form evidence that motivate the need for a structural learning model. Gamblers who incur a single loss return to the casino ten days later than gamblers who incur a single win, but the return time increases as more losses are incurred: gamblers with four past losses return about thirty days later than a gambler with four past wins. These findings suggest that gamblers incorporate outcomes from multiple past trips into the decision process. Reduced form evidence also shows gamblers' return times are significantly influenced by their beliefs about the house advantage and the uncertainty surrounding their beliefs.

Next, I estimate a structural model of the return time using a dynamic discrete choice framework. Structural methods allow for counterfactual predictions about how changes in marketing policies will affect consumer behavior (Reiss, 2011). In addition, structural estimation is appealing because it captures the dynamic forward-looking behavior of individuals. One obstacle to adopting structural estimation methods has been its computational burden, which is mainly due to two reasons. First, the likelihood is based on the explicit solution of a dynamic programming (DP) model. This requires us to obtain the fixed point of a Bellman operator for each possible point in the state space. Second, the number of points in the state space increases exponentially with the dimensionality of the state space, commonly known as the "curse of dimensionality". Imai et al. (2009) introduce a full-solution Bayesian approach to estimation of structural parameters. The key innovation in their algorithm is that they only need to conduct a single iteration of the Bellman operation during each estimation step (i.e., each MCMC draw). While conventional methods estimate the model only after solving the DP problem, their approach simultaneously solves the DP problem during parameter estimation. Because of this, the computational burden of their method is similar to that of non-Bayesian approaches but still intractable for dynamic learning models. In this paper I use forward simulation (see Bajari et al. (2007) and Hotz et al. (1994)). This significantly reduces the computation time and makes Bayesian estimation of a complex learning model feasible.

Learning models were first applied to marketing by Roberts and Urban (1988) and Eckstein et al. (1988). The initial models were relatively simple due to limitations on computer processing speeds and estimation algorithms available at the time. Erdem and Keane (1996) represents a significant methodological advance because it expanded the the class of learning models that became feasible to estimate. They used the method of Keane and Wolpin (1994) to obtain a fast and accurate approximate solution to the dynamic optimization problem and used simulation methods to approximate the likelihood function.

Recently, only a few research papers have applied the modified Bayesian MCMC algorithm first proposed by Imai et al. (2009) to estimate learning models with forward-looking consumers (see Roos et al. (2013)). Osborne (2011) is the first paper to allow for both learning and switching costs as sources of state dependence in a forward looking learning model. He also incorporated unobserved heterogeneity, however his paper assumed a "one-shot" learning model where only one purchase occurrence is all that is needed to learn everything about the product. While this may be reasonable in the product category he analyzed (laundry detergent) the purchase to purchase variability in most settings, including gambling, does not make this a suitable approach. The model proposed allows for learning to evolve over multiple exposures and also allows for individual parameters to be estimated using Bayesian methods with flexible mixture distributions on the heterogeneous parameters.

The results show that gamblers' prior beliefs overestimate the house advantage by a factor of about four and the counterfactuals suggest that this may be costing the casino substantial amounts of revenue. When gamblers overestimate the house advantage they overestimate projected future expenditures, which leads to delays in time until the next trip. The counterfactuals also illustrate the value of incorporating the beliefs and outcomes into the targeted marketing decision. I show that naive targeting strategies based on simple outcome heuristics are not sufficient and that more sophisticated targeting strategies can improve target marketing decisions. The final simulation shows that marketing strategies which vary offer amounts based on gamblers' beliefs in the house advantage improve profitability by close to 20%.

The Casino Industry

The gaming industry is a critical component of the U.S. economy. Casino gaming revenues are now at an historic peak at $67 billion. Oxford Economics (2014) estimates the U.S. gaming industry contributes nearly $240 billion to the national economy. With nearly 1,000 casinos operating in 39 states, Americans spend more money enjoying casino entertainment than they spend on spectator sports like football, baseball, basketball and soccer combined. Despite its substantial contribution to the economy, there are only a handful of papers that study effects of casino marketing (Nair et al., 2013; Park and Manchanda, 2015; Narayanan and Manchanda, 2012).

The casino industry has long understood the importance of effective customer relationship management (Compton, 1999). In today's gaming environment, a sophisticated tracking system is essential to remain competitive, especially in saturated markets such as Las Vegas and Atlantic City (Kilby and Fox, 1998). Casino marketing offers typically include a combination of free room nights (if the casino has a hotel) and slot promotional credits. Offers can also include additional complimentaries (or "comps") for virtually any other amenity available at the casino, such as show tickets, spa treatments, or dining credits.

To determine the optimal level of comps to offer their players, managers need to estimate gaming revenue at the player level. Casinos use player rating systems to track individual play and store player information. The marketing department uses this data to segment customers into tiers for offer generation.

Casinos record player information by having them enroll in the casino's loyalty program. Most casinos have enrollment centers on the casino floor. To incentivize gamblers to have their play tracked, casinos offer rewards programs. There are many differences between rewards programs across casinos, but the common denominator is that the magnitude of the reward is a function of gaming volume and possibly other on-property purchases. A rewards system increases the likelihood that the casino has the complete play history recorded because without rewards players are unlikely to allow the casino to track their play. While no-card play (that is, gambling without a loyalty card) still contributes a substantial portion to a casino's revenues, the amount of no-card play per person is often insignificant and unlikely to be of interest to the casino. Frequent gamblers often understand that it is in their best interest to have all of their play tracked in order to earn the highest reward possible.

Slot Machines

Electronic gaming machines are the most popular game among casino visitors, as more than half (51 percent) choose slot machines or video poker as their favorite game (American Gaming Association, 2013). For this and two other reasons I limit the analysis to gamblers that only play slot machines. First, tracking table games activity at the individual player level is still a very manual process and often inaccurate. On slot machines all play is recorded electronically through each gambler's loyalty card and because of this revenue is exact down to the penny. The second reason for limiting the analysis to slot players is because the skill of a tablegame player can to some degree dictate the outcome. For example, a skilled blackjack player can reduce the house advantage to nearly zero (or negative if they are counting cards) but an unskilled player can lose far more than expected in the long run. On slot machines it doesn't matter who is pulling the lever (or nowadays more often pressing the "spin" button); the outcome is completely random and in the long run the hold percentage should converge to the house advantage regardless of the individual gambler.

Before proceeding, there are a few industry terms used in this analysis that are common to the management of a slot department and will be important for the understanding of the model's learning process. *Handle* measures the total amount of money wagered on the machine. This measure of volume allows management to monitor the overall popularity of games. The *hold percentage* is the percentage of handle that that the machine keeps. This value fluctuates in the short term based on the randomness of the machine. The Las Vegas Strip hold percentage is around 7%, in Reno it is about 5%.[2] The *house advantage* is the long run expected hold percentage of the slot machine and depends on the payouts and odds specific to that machine. Slot machine advantages range from as low as .5% to as much as 25%. Many gaming jurisdictions have established minimum levels at which slot machines must pay back in order to prevent casino operators from placing too great a disadvantage on players (Kilby and Fox, 1998). The important point to remember is that the hold percentage reflects the *actual* amount of money kept by the casino, while the house advantage is what the casino *expects* to keep.

[2] http://gaming.unlv.edu/reports/nv_slot_hold.pdf

Theoretical and Actual Outcomes

Casino operators track both actual and theoretical player losses. The theoretical loss (also called "theo") is the amount of money the player was expected to lose. It is based on the following formula:

$$\begin{aligned} \text{Theoretical Loss} &= \text{Avg. Bet} \times \text{Decisions per Hour} \times \text{Hours} \times \text{House Advantage} \\ &= \text{Handle} \times \text{House Advantage} \end{aligned}$$

As mentioned earlier, for targeting purposes casinos typically ignore the actual outcomes and instead value their players on theoretical losses alone. The primary reason for doing so is to control for the randomness of outcomes. This creates a significant advantage for an analyst studying the impact of outcomes because marketing offer values are now almost completely randomly distributed across the actual outcomes. For example, if one gambler loses $500 and the another wins $4,000 but they were both *expected* to lose $300 they will receive the same offer.

Data

The dataset used to estimate the structural model comes from a large casino in the United States. The dataset includes the complete trip histories from over 28,000 randomly selected slot gamblers with around 110,000 trips occurring between February 2006 and May 2015. I observe basic demographic information such as gender, distance to the casino, age, and current loyalty card level (either "Silver" or "Gold").[3] A "trip" is defined as a distinct period of time where gambling activity is observed. For instance, a new trip is initiated when either 1) a player inserts their loyalty card into a slot machine for the first time or 2) a significant lapse in play occurs. The lapse required before a new trip is started is set by management to ensure that each trip record captures a distinct return to the casino rather than simply a suspension in play within a single trip. Typically the cutoff is three days, meaning that if no activity is observed after three days the trip is ended and any future play initiates a new trip.[4]

Each trip record includes detailed information about the gambling activity from that trip. The variables of interest are the start and end dates of the trip, actual and theoretical loss values, time played, average bet, promotional credits redeemed, comps received, and whether they stayed at the hotel or not. I also observe all marketing activity for these gamblers. Over the observed period, the gamblers redeemed over 2,500 separate

[3]Since I cannot reveal the actual labels used by the anonymous casino, I will refer to the upper loyalty tier as the "Gold" level.

[4]Note that after the three day lapse the trip is still recorded as "ending" on the last day of play, not three days later.

offers. For each offer I observe the period for which the offer is active (typically about 2.5 months), the date when the offer was sent to the gamblers (the "drop" date), the total number of promotional credits in the offer, and the comp room type. A slot promotional credit is essentially free slot play where any winnings from the promotional credit can be kept. The promotional value itself cannot be converted to cash.

Descriptive Analysis and Reduced Form Evidence

Before discussing the structural estimation procedure, I will first describe the data and show reduced form evidence that gambling outcomes impact the timing of the return trip. I exclude players whose play is high enough to warrant a casino host. This ensures that the only marketing communication between the casino and the gambler is done through direct marketing offers. I only keep gamblers who play slot machines exclusively for reasons discussed previously, which represents 37% of the low end player base. I remove gamblers whose first trip to the casino occurred before the first available marketing offer data is observed so that I have the complete marketing and trip history for each gambler. A few miscellaneous outliers are also removed due to bad data, more details are available upon request. Table 1 summarizes the cleaned data. For estimation, I only use gamblers with at least three trips to the casino. This is done to ensure that each gambler observes a sufficient amount of variance in the experienced hold percentages. In general, their statistics are quite similar to the aggregate level statistics.

INSERT TABLE 1 ABOUT HERE

Figure 1 shows the distribution of average uncensored intervisit times across gamblers with at least three trips at the casino. The median return time is about ten months.

INSERT FIGURE 1 ABOUT HERE

If gamblers learn from experience, their sensitivity to a single trip's gambling outcomes should decline over time. Experienced gamblers have more certainty about expected outcomes and therefore should be less likely to be swayed by their most recent outcome. Inexperienced gamblers (those with only a few trips) project expected outcomes only using a few signals, which can vary greatly between players and cause biases depending on the sequence realized. Over time gambler beliefs will converge to the truth and a single outcome will not have as much of an impact on the return decision. We see evidence for this in the data. Figure 2 shows the difference in median return times after a gambler lost compared with after a gambler won. The differences are grouped by experience, represented by the number of trips to the casino when the win or loss was realized. For example, when gamblers have less than five casino trips they tend to return about ten days later when they lose versus when they win. With more trips (and more experience) the difference in return times diminishes and gamblers become less impacted by the most recent trip's win or loss. To handle selection bias (that is, gamblers who eventually have many trips at the casino are inherently different from those with fewer trips) I only include players who eventually have between 15 and 30 total trips at the casino.

INSERT FIGURE 2 ABOUT HERE

To provide additional evidence for learning across trips, Figure 3 compares the difference in return times between winning and losing streaks, where the streak occurs over the past one through four trips. Ignoring any streaks (where the gambling streak equals one), players that lose tend to come back 10 days later than those that win. However, as players lose multiple times in a row, they delay the return trip by even greater lengths. For example, a gambler who lost the past three trips in a row will return about forty days later compared with a gambler who won the past three trips in a row. This suggests that the gamblers change their expectations as more outcomes are experienced.

INSERT FIGURE 3 ABOUT HERE

Next I provide reduced form evidence that gamblers' return times to the casino are influenced by their beliefs in the house advantage. I estimate return times using a weibull hazard model and include the posterior mean of the house advantage and its posterior variance as covariates - the posteriors are generated from a Bayesian learning process (the updating process is discussed in detail later). I also include a variety of demographic and last trip variables: age, sex, card level, distance to the casino, whether they stayed at the hotel on the last trip, whether they redeemed a promotion on the last trip, the log of the total comps received, trip length, and last trip theoretical loss.

In a Bayesian learning process, the prior mean and prior variance dictate the evolution of the posterior mean and variance. Because of this, in a reduced form hazard model a prior mean and variance needs to be selected in order to generate the posterior beliefs on the house advantage (the mean and its uncertainty). In the structural model these priors are estimated, but for reduced form evidence I estimate 150 hazard models over a grid of 15 prior mean and 10 prior variance starting points. Figure 4 shows examples of the truncated normal shapes to illustrate the variety of prior settings that are considered for the house advantage beliefs. The idea is that by estimating many hazard models over this grid I can determine if the reduced form coefficient estimates are sensitive to the learning process priors. The specific gridpoints are available in the Web Appendix.

INSERT FIGURE 4 ABOUT HERE

Figure 5 plots the coefficient on the posterior mean across all 150 gridpoints. The latticed plane is positioned where the z-axis equals zero; any points above this plane are positive and below are negative and points that are filled in are significant at the .05 level. Except for very low prior variance values (where convergence of weibull is not achieved) the coefficients on the posterior mean tend to be positive and significant, which means that the return time increases with the posterior mean of the house advantage. In other words, as the belief in the house advantage increases gamblers take longer to return. Figure 6 shows a similar plot but for the coefficient on the posterior variance. At very low prior variance settings the coefficient cannot be estimated. The coefficients that can be estimated are significant, suggesting that uncertainty in the gamblers' beliefs on the house advantage influence the return decision.

The reduced form evidence suggests that 1) learning should be incorporated into a model of the return times, and 2) posterior beliefs in the house advantage (the mean and its uncertainty) influence the return time. The drawback of a reduced form approach is that it does not account for any forward-looking behavior of the gamblers.

INSERT FIGURE 5 ABOUT HERE

INSERT FIGURE 6 ABOUT HERE

A Model of Gambler Learning

In this section I propose a structural model of the casino return decision process. I first outline the dynamic optimization problem somewhat generally and then introduce the learning component and specific utility function.

Gambler Dynamic Optimization Problem

I model casino return times in the framework of a dynamic discrete choice model, which can be interpreted as a generalization of a structural hazard model Van den Berg (2001). I estimate an infinite horizon model of a forward looking agent. Each decision period the gambler decides to return to the casino or not by comparing their current and discounted future utilities of each action.

Let θ be the J-dimensional parameter vector. Let S be the finite set of state space points and s be an element of S. Let A be the finite set of all possible actions and a be an element of A. Let $u(s, a, \varepsilon_a, \theta)$ be the current period utility of choosing action a given state s and ϵ is a vector whose ath element is a random choice to the current returns of choice a. The transition probability of next period state s', given current state s and action a is $f(s'|s,a,\theta)$. Given a discount rate β, The time invariant value function can be defined to be the maximum of the discounted sum of expected utilities:

$$V(s_t,\varepsilon_t,\theta) \equiv \max_{\{a_t,a_{t+1},\ldots\}} \mathbb{E}\left[\sum_{\tau=t}^{\infty} \beta^{\tau} u(s_\tau,a_\tau,\epsilon_{a_\tau},\theta)\,|s_t,\varepsilon_t\right]$$

This value function is known to be the unique solution to the Bellman equation

$$V(s,\varepsilon,\theta) = \max_{a\in A}\left\{\mathbb{E}\left[u(s,a,\varepsilon_a,\theta)\right] + \beta\mathbb{E}_{s',\varepsilon'}\left[V(s',\varepsilon',\theta)\,|s,a\right]\right\}$$

The first expectation is included because even when making the decision to return to the casino the utility is not known until after the trip has been realized. The second expectation is taken with respect to the next period shock ε' and the next period state s'.

If we define $EV(s,a,\varepsilon_a,\theta)$ to be the expected value of choosing action a then

$$EV(s,a,\varepsilon_a,\theta) = \mathbb{E}\left[u(s,a,\varepsilon_a,\theta)\right] + \beta\mathbb{E}_{s',\varepsilon'}\left[V(s',\varepsilon',\theta)\,|s,a\right]$$

and the value function can be written as

$$V(s,\varepsilon,\theta) = \max_{a\in A} EV(s,a,\varepsilon_a,\theta)$$

The dataset for estimation includes variables which correspond to state vector s and choice a but the choice shock ε is not observed. I observe data for $i = 1,\ldots,N$ gamblers, and each gambler i has T_i observations. The observed data for individual i is denoted $y_i^d \equiv \left\{a_{i,t}^d, s_{i,t}^d\right\}_{t=1}^{T_i}$ and $Y^d \equiv \left\{y_i^d\right\}_{i=1}^{N}$ with superscript d to represent that this is observable data. Furthermore,

$$a_{i,t}^d = \arg\max_{a\in A} EV\left(s_{i,t}^d, a, \varepsilon_a, \theta\right)$$

Let $\pi(\cdot)$ be the prior distribution of θ and let $L\left(Y^d|\theta\right)$ be the likelihood of the model, given the parameter θ and the value function $V(\cdot,\cdot,\theta)$, which is the solution of the dynamic programming problem. Then we have a posterior distribution function of θ:

$$P\left(\theta|Y^d\right) \propto L\left(Y^d|\theta\right)\pi(\theta)$$

Let $\varepsilon_i \equiv \{\varepsilon_{i,t}\}_{t=1}^{T_i}$ and $\varepsilon \equiv \{\varepsilon_i\}_{i=1}^{N}$. Because ε is not observed to the analyst, the likelihood is an integral over it. That is, if we define $L\left(Y^d|\varepsilon,\theta\right)$ to be the likelihood conditional on (ε,θ), then

$$L\left(Y^d|\theta\right) = \int L\left(Y^d|\varepsilon,\theta\right) dF_{\varepsilon}\left(\varepsilon|\theta\right)$$

The value function enters into the likelihood through the choice probability. Per-period utility is specified as follows:

$$u\left(s,a,\varepsilon_a,\theta\right) = \widehat{u}\left(s,a,\theta\right) + \varepsilon_a$$

Where $\widehat{u}\left(s,a,\theta\right)$ is the deterministic component of the per-period utility. Also,

$$\widehat{EV}\left(s,a,\theta\right) = \mathbb{E}\left[\widehat{u}\left(s,a,\theta\right)\right] + \beta \sum_{s'} \mathbb{E}_{\varepsilon'}\left[V\left(s',\varepsilon',\theta\right)\right] f\left(s'|s,a,\theta\right)$$

Then, generally we have:

$$\Pr\left[a^d_{i,t}|s^d_{i,t},V,\theta\right] = \Pr\left[\varepsilon_a - \varepsilon_{a^d_{i,t}} \leq \widehat{EV}\left(s,a^d_{i,t},\theta\right) - \widehat{EV}\left(s,a,\theta\right); \forall a \neq a^d_{i,t}|s^d_{i,t},V,\theta\right]$$

I assume that each ε_a is distributed independent, identically extreme value. In addition, I introduce a hierarchical structure so that each V and θ are specific to gambler i. In the empirical application since A contains two actions, either return to the casino $(a=1)$ or not $(a=0)$, the conditional choice probabilities take the following form:

$$\Pr\left[a^d_{i,t}=1|s^d_{i,t},V_i,\theta_i\right] = \frac{1}{1+\exp\left(-\left[\widehat{EV}\left(s^d_{i,t},1,\theta_i\right) - \widehat{EV}\left(s^d_{i,t},0,\theta_i\right)\right]\right)}$$

Introducing a Hierarchical Structure

To allow for parameter estimates to vary by individual characteristics, I introduce a hierarchical structure. Understanding individual differences is crucial in strategic CRM applications when developing targeted marketing strategies (Rossi et al., 2005). The hierarchical parameters are specified as a function of an individual's observable characteristics. We have nz observable characteristics on each individual. If we let Z denote a matrix with N rows and nz columns and similarly Θ be a matrix of N rows and J columns, where the ith row of Θ is the parameter estimates for individual i then we have:

$$\Theta = Z\Delta + U$$

Where Δ is a $nz \times J$ matrix of coefficients on the observables and U is a vector of residuals. This is simply a multivariate regression of Θ on Z. In each row of U,

$$u_i \sim N(0, \Sigma_\theta)$$

The priors are specified as follows:

$$\begin{aligned} \text{vec}(\Delta|\Sigma_\theta) &\sim N\left(\text{vec}\left(\bar{\Delta}\right), \Sigma_\theta \otimes A^{-1}\right) \\ \Sigma_\theta &\sim IW(\nu, \Sigma) \end{aligned}$$

Hierarchical models for panel data structures are ideally suited for MCMC methods. A Gibbs-style Markov chain can be constructed by considering the two sets of conditionals:

$$\begin{aligned} &\theta_i|\tau, y_i^d \\ &\tau|\{\theta_i\} \end{aligned}$$

The first line exploits the fact that the θ_i are independent, conditional on the first stage priors $\tau = \{\Delta, \Sigma_\theta\}$. The second line exploits the fact that $\{\theta_i\}$are sufficient for τ. That is, once the individual level parameters are drawn they serve as "data" to the inferences on the priors. Due to the non-linearity of the model proposed, there is no way to conveniently sample from the conditional posterior (i.e., using a Gibbs sampler). For this reason, I employ a Metropolis algorithm to draw θ_i For each gambler i, I draw candidate random effects parameters θ_i^n by perturbing the current draw θ_i^o: $\theta_i^n = \theta_i^o + \varepsilon$, where $\varepsilon \sim \mathcal{N}\left(0, s^2\Sigma\right)$. I then compare the likelihood of the new parameters with the old parameters and accept the new parameters with probability α:

$$\begin{aligned} \alpha(\theta_i^n, \theta_i^o) &= \min\left\{1, \frac{L_i\left(Y_i^d|\theta_i^n\right) q\left(\theta_i^n, \theta_i^o\right)}{L_i\left(Y_i^d|\theta_i^o\right) q\left(\theta_i^o, \theta_i^n\right)}\right\} \\ &= \min\left\{1, \frac{L_i\left(Y_i^d|\theta_i^n\right)}{L_i\left(Y_i^d|\theta_i^o\right)}\right\} \end{aligned}$$

The second line is a result of the symmetry of the transition density $q(\cdot,\cdot)$.

Learning About the House Advantage

In this section I introduce the learning process. As gamblers play slot machines, they receive signals about that casino's house advantage. Before receiving any information, they have a truncated normal prior belief on the house advantage:

$$A_i \sim \mathcal{TN}\left(A_{0i}, \sigma_{0i}^2, 0, 1\right)$$

In other words, before a gambler's first trip to the casino, they expect to lose a certain percentage of every dollar cycled through the slot machine. The house advantage is bounded from below at zero because it is irrational for a gambler to expect to win money from a slot machine in the long run. It is also bounded from above at one because in the long run it is impossible for a machine to pay out more money than is put into it. Again, in the short term the *hold percentage* can fall outside of these bounds, but the gambler's beliefs on the *house advantage* cannot reasonably be outside of this range.

The player's experience at the casino does not fully reveal the house advantage because of the inherent variability of gambling outcomes. As previously discussed, there is quite often a difference between the *hold percentage* and the *house advantage* for gambler i on occasion t. I denote the hold percentage as H_{it}, which can be interpreted as the "experienced" house advantage, and the house advantage as A_i. The hold percentage is thus a noisy signal of the house advantage:

$$H_{it} = A_{it} + \eta_{it}, \text{ where } \eta_{it} \sim \mathcal{N}\left(0, \sigma_\eta^2\right)$$

The gamblers update their posterior mean and variance of the house advantages using a Bayesian updating process. That is:

$$A_{it} = \frac{\sigma_{0i}^2}{N_i\left(t\right) \cdot \sigma_{0i}^2 + \sigma_{\eta_i}^2} \sum_{s=1}^{t} H_{is} d_{is} + \frac{\sigma_{\eta_i}^2}{N_i\left(t\right) \cdot \sigma_{0i}^2 + \sigma_{\eta_i}^2} A_{0i}$$

$$\sigma_{it}^2 = \frac{1}{1/\sigma_{0i}^2 + N\left(t\right) \cdot 1/\sigma_{\eta_i}^2}$$

where $N_i\left(t\right)$ is the number of gambling experiences realized up through time t and d_{it} is an indicator for whether the player gambled at time t. The Appendix contains a proof showing that if the prior is truncated normal and the signal is an unbounded normal then the corresponding posterior is also a truncated normal.

Figure 7 plots the distribution of the hold percentage and house advantages across all trips in the dataset. The hold percentage distribution is what the gamblers experience and the house advantage is what gamblers are attempting to learn about. The dashed vertical line is the mean house advantage: with enough exposures the gamblers will learn this value with certainty if slot machines are selected at random. Notice there is a small bump for high payoff (i.e., low house advantage) games. A more complex learning model would account for this multi-modality, but in this application learning about the mean alone seems reasonable given that the distribution it so concentrated. Note that each gambler's experienced house advantage is observed by the analyst, so even if gamblers don't select slot machines at random (e.g., they only play one machine that happens to have a very low house advantage) the analyst can still determine if their estimated posterior beliefs are above or below the true house advantage. However, since the distribution of the house advantages is so tightly centered for simplicity I assume that the machines are selected at random and only the mean house advantage matters.

INSERT FIGURE 7 ABOUT HERE

Cost of Gambling

When gamblers consider a return trip to the casino they need to form projections on the cost of gambling. This influences the expected future utilities. Under perfect knowledge expected cost is the same as the theoretical loss:

$$\text{Gambling Cost} = \text{Avg. Bet} \times \text{Decisions per Hour} \times \text{Hours} \times \text{House Advantage}$$

However, since gamblers have imperfect knowledge on the house advantages there is uncertainty in projections of their gambling costs. This uncertainty depends on their current beliefs at time t:

$$\text{Gambling Cost} \sim \mathcal{N}\left(\text{BDH} \cdot A_{it}, \text{BDH}^2 \cdot \sigma_{it}^2\right)$$

"BDH" represents the product of average bet, decisions per hour, and hours played. These three variables are completely within the gambler's control (i.e., there is no uncertainty) and represent the gambler's play style.

It is important to note that the gambler's projected average bet, game speed, and time may be a function of their current beliefs on the house advantage. For example, if faced with a relatively high house advantage players may decide to decrease their average bet to reduce projected gambling costs (everything else held constant). Similarly, higher uncertainty in the cost may lead to play that is more likely to result in a lower cost. Furthermore, gamblers may also adjust their play style based on currently available marketing offers. For instance, a gambler returning to the casino on a free room offer may play more aggressively than usual since the comped room frees up money that could be used for gambling. To account for this, the BDH value can vary during the forward simulation (discussed in more detail later).

Utility Specification

In this section I introduce the utility function. The utility associated with returning to the casino is given by the following expression:

$$
\begin{aligned}
u(a=1,s,\varepsilon_1,\theta)_{it} &= \theta_{1i}\left(\mathrm{BDH}_{it}\cdot H_{it}\right)+\theta_{1i}r_i\left(\mathrm{BDH}_{it}\cdot H_{it}\right)^2 \\
&\quad +\theta_{2i}\text{Offer Gaming Value}_{it} \\
&\quad +\theta_{3i}\text{Offer Room Value}_{it} \\
&\quad +\Omega f\left(w_{it}\right)+\Gamma\text{Month}_{it}+\theta_{0i}+\varepsilon_{1it} \\
u(a=0,s,\varepsilon_0,\theta)_{it} &= \varepsilon_{0it}
\end{aligned}
$$

Where u_{it} is the utility for gambler i at time t. BDH is the product of average bet, pulls per hour, and hours played.[5] BDH multiplied by the hold H it captures the gambling expense realized from that trip. Importantly, this expense is not known at the time of the decision and only realized after experiencing the outcome. θ_1 represents the utility weight gamblers attach to this cost, r is the risk coefficient, θ_2 is the utility weight of the offer's gaming value, and θ_3 is the utility weight of the offer's room value. Ω is the vector of utility weights associated with a function of the time since the last trip (w), which I specify as polynomials: $\Omega f(w)=\omega_1 w+\omega_2 w^2+\omega_3 w^3+\omega_4 w^4+\omega_5 w^5$. Γ is the vector of utility weights associated with the month the decision to return was made, in order to capture impacts from seasonality: $\Gamma\text{Month}=\gamma_1\mathbb{I}\left[\text{Month}=1\right]+\ldots+\gamma_{11}\mathbb{I}\left[\text{Month}=11\right]$. θ_0 is an intercept. ε is the random component associated with this choice, which is known to the gambler but not observed by the analyst. Ω and Γ are common across individuals, while θ_0, θ_1, θ_2, and r are specific to the individual.

Given the utility specification and the learning process, expected utility is given by the following:

$$
\begin{aligned}
\mathbb{E}_{A_{it}}\left[u(a=1,s,\varepsilon_1,\theta)\right] &= \theta_{1i}\left(\mathrm{BDH}_{it}\cdot A_{it}\right)+\theta_{1i}r_i\left(\mathrm{BDH}_{it}\cdot A_{it}\right)^2+\theta_{1i}r_i\left(\mathrm{BDH}_{it}^2\cdot\left(\sigma_{it}^2+\sigma_{\eta_i}^2\right)\right) \\
&\quad +\theta_{2i}\text{Offer Gaming Value}_{it} \\
&\quad +\theta_{3i}\text{Offer Room Value}_{it} \\
&\quad +\Omega f\left(w_{it}\right)+\Gamma\text{Month}_{it}+\theta_{0i}+\varepsilon_{1it} \\
u(a=0,s,\varepsilon_0,\theta) &= \varepsilon_{0it}
\end{aligned}
$$

Under this specification, utility is linear in the cost of gambling. As in Erdem and Keane (1996), the formulation is such that given a strictly negative θ_1, utility is concave in A for $r>0$, linear in A for $r=0$, and convex for $r<0$. Thus if there is uncertainty about the house advantage, the consumer is risk averse, risk neutral or risk seeking as $r>0$, $r=0$, or $r<0$, respectively. As noted earlier, the uncertainty is in the beliefs on the house advantage, even in the "current" decision period. Furthermore, while the offer values are known in the current period they are not known in future periods, so gamblers form expectations over these

[5] Note BDH is the same as the handle, but in order to prevent confusion between handle and hold I call it BDH

values as well. In the simulation I draw values from the empirical joint distribution of room and gaming offer values.

Model Estimation

The Estimation Procedure

The structural parameters of interest are $\{\theta_{0i}, \theta_{1i}, \theta_{2i}, \theta_{3i}, r_i, \Omega, \Gamma\}$ and the priors on each individual's learning process $\{A_{0i}, \sigma_{0i}^2\}$. The proposed estimation procedure uses the advantages of Bayesian estimation (versus classical estimation methods) while remaining computationally feasible. The biggest challenge presented when estimating structural learning models is that the state space is incredibly large. When discounting future expectations, a forward looking gambler needs to consider the impacts of all potential outcomes and the associated implications on the learning process itself. For example, the specific hold percentage a gambler expects to experience on a return trip will influence how their posterior beliefs update, which in turn influences later return decisions. Clearly, evaluating every single potential learning path is daunting and because of this a full-solution Bayesian approach is not feasible, such as the method proposed by Imai et al. (2009).

Erdem and Keane (1996) use backwards induction to solve their learning model. However, the entire backwards induction needs to be re-solved at every parameter estimate. This is not feasible for Bayesian methods, which typically rely on tens of thousands of MCMC draws to converge onto the posteriors. The impracticality of their method is not limited simply to the desire to use Bayesian rather than classical methods: the complexity of the proposed utility function and the hierarchical structure also render their approach as unfeasible.

Rather than attempt to visit every single learning path I forward simulate over many potential paths and discount the simulated values. More likely paths will be simulated more often and averaging over many simulated paths provides a consistent estimate of the discounted future returns. The advantage of this approach is that if the utility function is linear in parameters we only have to simulate the paths once for each considered starting state since the current parameter estimates do not affect the simulated values (see Hotz et al. (1994) for a discussion of this method). The discounted terms can be separated from the parameters such that the parameters simply scale the discounted values during estimation.

One challenge is that in the utility function specified some variables enter non-linearly, namely the prior mean and prior variance of the beliefs in the house advantage. Note that the variability in the hold percentage ($\sigma_{\eta_i}^2$) is observed by both the analyst and the gambler, so there is no need to estimate this. To handle the non-linearity of the learning priors, I forward simulate over a grid of prior mean and prior variance values and during likelihood evaluation use bi-linear interpolation to fill in areas near the simulated prior mean and variance gridpoints. The intuition is that the observed data should reflect a specific learning process with a particular prior mean and prior variance and during the MCMC iterations I search over the prior learning parameters that maximize the likelihood. One disadvantage of the estimation approach is that the discount factor cannot be estimated and needs to be selected prior to the forward simulation procedure. However, the estimation strategy makes it relatively easy to compare a few candidate discount factors by simply adding additional parallelized forward simulations. More details on the forward simulation algorithm and bi-linear interpolation are available in the Appendix and the Web Appendix.

After forward simulation is complete the MCMC draws can proceed at usual speed. At first glance, it appears that we have simply pushed the computational intractability to the front of the estimation process, but it is important to note that each forward simulation for each grid point and each starting state can be run at the same time. With enough computers the whole procedure can be completed in minutes due to its massively parallel nature. Once the forward simulation is complete the discounted expected values are simply plugged into the likelihood and Bayesian estimation proceeds as usual.

Play Style Estimation

Since a portion of the gambling cost is within the player's control (average bet, decisions per hour, and time, or "BDH") I account for potential adjustments in a gambler's play style due during the forward simulation. For instance, if a gambler believes that the house advantage is very high they may decrease their next trip's average bet to reduce expected gambling costs. To estimate these effects, I estimate a regression of the log ratio of the next return trip's BDH relative to the previous trip's BDH:

$$\begin{aligned}\ln\left(\frac{\mathrm{BDH}_{i,t+1}}{\mathrm{BDH}_{i,t}}\right) &= \beta_1 A_{it} + \beta_2 \sigma_{it}^2 + \beta_3 g_{it} + \beta_4 g_{i,t+1} \\ &\quad + \beta_5 r_{it} + \beta_6 r_{i,t+1} \\ &\quad + \beta_7 o_{it} + \beta_0 + \varepsilon_{it}\end{aligned}$$

where g is the offer gaming value, r is the offer room value, and o is the outcome, represented as the casino's revenue from the player (positive values indicate a player loss), and$\varepsilon_{it} \sim \mathcal{N}\left(0, \sigma_\varepsilon^2\right)$. The coefficients on promotional credits and comp values control for any changes in play behavior attributed to reductions in the overall trip cost. For instance, if a player is returning on a free room offer they may increase their BDH.

It is important to note that these coefficients are in regards to the *play style*, not the return decision. For example, if the coefficient on house advantage is positive it simply means that *when* the player returns they tend to play more aggressively - it does not imply that higher gambling costs increase the utility of returning to the casino.

For each of the 150 prior mean and prior variance gridpoint combinations, I run 10,000 MCMC iterations (keeping only every 10th draw) and save the posterior means. The posterior means are used during the forward simulation for adjusting BDH values as more experience signals are realized. The priors are specified as follows:

$$\begin{aligned} \beta &\sim \mathcal{N}\left(\bar{\beta}, \sigma_{\varepsilon}^{2} \cdot A^{-1}\right) \\ \sigma_{\varepsilon}^{2} &\sim \left(\nu \cdot \text{ssq}\right)/\chi_{\nu}^{2} \end{aligned}$$

The coefficient estimates and prior settings are available in the Web Appendix.

Policy Function

In this section I outline the policy function used in forward simulation. The policy function estimates the probability of return given the current state. I use a Bayesian non-parametric method as outlined by Rossi (2014) to estimate this policy function. Non-parametrically, a regression models the conditional distribution of y given x. A fully non-parametric approach to regression uses the entire conditional distribution of y given x as the object of interest for inference. For the policy function I model the joint distribution of y and x and then use this joint distribution to compute the conditional distribution of $y|x$. This approach does not require assumptions and specific functional forms for how the x variables influence the conditional distribution of y.

For the policy function I estimate a five component mixture model. The covariates in x are the posterior mean and variance, predicted next trip BDH, the gambler's weeks since the last trip, month, and the room and slot promotional credit values if an offer is available during that week. I first approximate the joint distribution and then use these draws to compute the implied conditional distribution. Formally, for the rth draw with K mixing components:

$$f(y,x)^r = \sum_{k=1}^K \pi_k^r \phi\left(y,x|\mu_k^r,\Sigma_k^r\right)$$
$$f(y|x)^r = \frac{\sum_{k=1}^K \pi_k^r \phi\left(y,x|\mu_k^r,\Sigma_k^r\right)}{f(x)^r}$$
$$f(x)^r = \int f(y,x)^r \, dy = \sum_{k=1}^K \pi_k^r \bar{\phi}_k(x)^r$$
$$\bar{\phi}_k(x)^r = \int \phi\left(y,x|\mu_k^r,\Sigma_k^r\right) dy$$

I use a finite mixture of normals model to simulate from the joint posterior density. The mixture of normals model is written as follows:

$$y_i \sim N\left(\mu_{\text{ind}i}, \Sigma_{\text{ind}i}\right)$$
$$\text{ind}_i \sim \text{Multinomial}(\pi)$$

Here y_i is a two dimensional vector and π is a vector of K mixture probabilities. Priors for the model are specified in conditionally conjugate forms:

$$\pi \sim \text{Dirichlet}(\alpha)$$
$$\mu_k \sim N\left(\bar{\mu}, \Sigma_k \otimes a_\mu^{-1}\right), \quad k = 1,\dots,K$$
$$\Sigma_k \sim IW(\nu, V)$$

Any functional of the conditional distribution such as the conditional mean can be computed based on the rth draw of the joint distribution. In the policy function, I use the conditional mean in the policy regression. The linear structure of the mixture of normals model can be exploited to facilitate computation of the conditional mean. More details are available in the Web Appendix.

$$\begin{aligned}
\mathbb{E}[y|x] &= \int y f(y|x)\, dy = \int y \frac{\sum_k \pi_k \phi\left(y,x|\mu_k,\Sigma_k\right)}{f(x)} dy \\
&= \frac{1}{f(x)} \int y \sum_{k=1}^K \pi_k \phi_k(y,x)\, dy \\
&= \frac{1}{f(x)} \sum_{k=1}^K \pi_k \int y \phi_k(y,x)\, dy \\
&= \frac{1}{f(x)} \sum_{k=1}^K \pi_k \int y \frac{\phi_k(y,x)}{\bar{\phi}_k(x)} \bar{\phi}_k(x)\, dy \\
&= \frac{1}{f(x)} \sum_{k=1}^K \pi_k \mathbb{E}_k[y|x]
\end{aligned}$$

Estimating the Value Function

Suppose $\sigma(s,\varepsilon)$ is the optimal action given state s and shock ε based on the policy function estimated in the previous section. Following Bajari et al. (2007), I take advantage of the fact that for a given learning process prior mean and prior variance, the parameters enter the utility linearly.

$$
\begin{aligned}
\mathbb{E}\left[u\left(a=1,s,\varepsilon_1,\theta\right)\right] &= \Big[\left(\text{BDH}\cdot A_{it}\right) \quad \left(\text{BDH}\cdot A_{it}\right)^2 \quad \left(\text{BDH}^2\cdot\left(\sigma_{it}^2+\sigma_{\eta_i}^2\right)\right) \\
&\qquad \text{Gaming Value}_{it} \quad \text{Room Value}_{it} \quad f\left(w_{it}\right) \quad \text{Month}_{it} \quad 1\Big]\cdot \\
&\quad \left[\theta_{1i} \quad \theta_{1i}r_i \quad \theta_{2i} \quad \theta_{3i} \quad \Omega \quad \Gamma \quad \theta_{0i}\right]'+\varepsilon_1 \\
&\equiv \Psi_{it}\cdot\left[\theta_{1i} \quad \theta_{1i}r_i \quad \theta_{2i} \quad \theta_{3i} \quad \Omega \quad \Gamma \quad \theta_{0i}\right]'+\varepsilon_1
\end{aligned}
$$

Defining:

$$
W\left(s;\sigma\left(s,\varepsilon\right)\right)=\mathbb{E}\left[\sum_{t=0}^{\infty}\beta^t\left(\sigma\left(s_t,\varepsilon_t\right)\Psi_{it}\right)|s_0=s\right]
$$

We then have

$$
V\left(s,\sigma;\theta\right)=W\left(s;\sigma\left(s,\varepsilon\right)\right)\cdot\theta
$$

Exploiting this allows us to forward simulate the data only once (for each prior mean and variance gridpoint). This eases the computational burden significantly, allowing us to use the stored values when searching over the θ parameters during MCMC draws.

Forward Simulation & Parallelization

By taking advantage of the massively parallel structure of the forward simulation the expected value terms can be computed in a manner of hours with a reasonably sized dataset so long as many processors are available. With recent advances in online computing, estimating this complex model becomes a relatively inexpensive and fast process. To execute the forward simulation process, I use Amazon's EC2 service which rents processors at an hourly rate.[6] With small memory loads the cost is very low (less than a penny per hour) so running hundreds of instances simultaneously for a few hours is quite inexpensive.

I simply use each record in the data directly as starting states because creating starting states intended to "cover" the state space of the data is just as complex and would also require interpolation. Note that each record represents one decision period (one week), so there are hundreds of records per gambler. To give some context as to the scale of the parallelization in the empirical estimation, I conduct 100 forward simulations for each decision at each of the 287,205 rows of data over each of the 150 prior mean and variance gridpoints. This implies that theoretically the process can be divided across 8,616,150,000 servers and completed nearly instantaneously. In reality I divide the process over 30,000 servers and the process is done in about 50 hours (the servers are not all initiated at once).

[6] https://aws.amazon.com/

At the end of the forward simulation, I obtain the expected values for returning or not at each record for each learning process prior mean and prior variance gridpoint. To recover the structural parameters, the three dimensional array (rows of data x discounted basis functions x 150 gridpoints) is then referenced during the MCMC process. I allow proposed prior mean and prior learning variances to take on any value within the range of gridpoints and use bi-linear interpolation to estimate the missing expected value. See the Web Appendix for more details on the bi-linear interpolation process.

Recap of the Estimation Procedure

For clarity, in this section I summarize the estimation procedure. First, I estimate the play style regression coefficients and policy function mixture components for each of the 150 prior mean and variance learning process parameter combinations. I do this because each prior mean and variance determines the evolution of the Bayesian updating process that each player experiences. Next, at each starting state and for each of the 150 learning process prior gridpoints I forward simulate using the play style regression coefficients and policy function parameters specific to that gridpoint. This process mimics a gambler projecting potential outcomes and then discounting the values that results in decision to return or not in a particular week. Since this process can be run in parallel across starting states and learning prior grid I divide the estimation over many cloud computers using Amazon EC2. Once the average discounted values are obtained for each of the 150 gridpoints and each record of the data, I then use standard Bayesian MCMC methods to estimate the structural parameters. As previously noted, the coefficients simply scale the values obtained from the forward simulation and because of this it is easy to introduce a hierarchical structure. The MCMC routine then searches for the structural parameters that make the observed data most likely. More details on the entire estimation procedure are available in the Web Appendix.

Identification

The structural parameters of interest are $\{\theta_{0i}, \theta_{1i}, \theta_{2i}, r_i, \Omega, \Gamma\}$ and the priors on each individual's learning process $\left\{A_{0i}, \sigma_{0i}^2\right\}$. Recall the choice specific value functions are as follows:

$$
\begin{aligned}
EV\left(s, a, \varepsilon_a, \theta\right) &= \mathbb{E}\left[u\left(s, a, \varepsilon_a, \theta\right)\right] + \beta \mathbb{E}_{s', \varepsilon'}\left[V\left(s', \varepsilon', \theta\right) | s, a\right] \\
EV\left(s, a=1, \varepsilon_1, \theta\right) &= \theta_{1i}\left(\mathrm{BDH}_{it} \cdot A_{it}\right) + \theta_{1i} r_i \left(\mathrm{BDH}_{it} \cdot A_{it}\right)^2 + \theta_{1i} r_i \left(\mathrm{BDH}_{it}^2 \cdot \left(\sigma_{it}^2 + \sigma_\eta^2\right)\right) \\
&\quad + \theta_{2i} \text{Offer Gaming Value}_{it} + \theta_{3i} \text{Offer Room Value}
\end{aligned}
$$

$$+\Omega f\left(w_{it}\right)+\Gamma\text{Month}_{it}+\theta_{0i}+\varepsilon_{1it}$$

$$+\beta\mathbb{E}\left[V\left(s',\varepsilon',\theta\right)|s,1\right]$$

$$EV\left(s,a=0,\varepsilon_{0},\theta\right) = \varepsilon_{0it}+\beta\mathbb{E}\left[V\left(s',\varepsilon',\theta\right)|s,0\right]$$

Suppose that gamblers had complete information about the casino's house advantage. This would imply that $A_{it}=A_i$ and $\sigma_{it}^2=0$, and results in us being unable to separately identify A_{it} and θ_{1i}. Since gamblers observe the variation in the hold percentages, σ_η^2 does not need to be estimated, unlike in Erdem and Keane (1996). Because the variability in hold percentages changes over time, it appears we can identify r_i. But since we cannot identify θ_{1i} in this complete information scenario, only the product $\theta_{1i}r_i$ is identified. So identification rests upon the assumption that incomplete information exists (which is true for a static model as well).

With incomplete information, the gambler's priors and their hold percentage exposures will guide the learning process path. Identifying the prior mean separately from the prior variance is challenging in most applications, the common solution being to fix the prior variance at one and estimate the signal variance and prior mean. But since I observe the signal variance I use the functional form of the Bayesian learning process to enable identification. A similar argument is made in Sriram et al. (2015). The priors determines how A_{it} and σ_{it}^2 evolve. Thus these parameters are pinned down by the extent to which new hold percentage signals change the probability of returning (and hence the actual returns observed in the data). The hold percentage exposures vary across gamblers and create variation in the evolution of A_{it} and σ_{it}^2. So even if every gambler started with the same learning priors, the variability in outcomes across gamblers allows us to identify r_i.

Results

The results are estimated on a random subsample of 1,000 gamblers. For each gambler, 100 paths were forward simulated to derive the discounted values.[7] To assist with parameter convergence in the hierarchical model, I first estimate a homogeneous model and use those parameters as the starting values in the hierarchical estimation. The parameters are estimated using a random-walk step on each MCMC draw. Since the parameter space is quite large, I partition the estimation into four parameter blocks to make the parameter search easier (Chib and Greenberg, 1996). The first block contains the learning process prior mean and variance $\left\{A_0,\sigma_0^2\right\}$, the second contains the cost, risk aversion, offer coefficients, and intercept$\{\theta_0,\theta_1,\theta_2,\theta_3,r\}$,

[7]The discounted values began to converge after averaging 50 forward simulated paths. I selected 100 to ensure consistency in the estimates.

the third block are the coefficients on the weeks since last trip polynomials $\{\Omega\}$, and the fourth block are the month control variables $\{\Gamma\}$. Details on the estimation procedure is available in the Web Appendix.

In the homogeneous model, I run 80,000 MCMC draws. I discard the first 60,000 draws and keep only every 10th draw thereafter. I initialized the chain using MLE estimates. The acceptance rates of each of the four blocks is between 15% and 50% and the likelihood is -15,950. Table 2 contains the posterior means of the kept draws. As expected the coefficient on the gambling expense is negative.

INSERT TABLE 2 ABOUT HERE

The homogeneous results are used as starting parameters for the hierarchical model. In the hierarchical model, I allow the learning process prior mean and variance, intercept, cost, risk coefficient, and offer coefficients to be a function of individual level information. The coefficients on the weeks since last trip polynomials $\{\Omega\}$and the month control variables $\{\Gamma\}$ remain fixed across the gamblers. The individual level covariates are the gambler's age, sex, distance to the casino, and an indicator for whether the gambler is at the "Gold" loyalty card status. I run 80,000 MCMC draws, discarding the first 60,000 and keeping every 10th draw thereafter. The model's likelihood is -11,358. This is a significant improvement over the homogeneous model and also greater than the likelihood from the same model with no forward looking (-11,401 when the discount factor $\beta = 0$ versus $\beta = .98$). Details on other model parameters are available in the Web Appendix.

Table 3 displays the estimates for the hierarchical parameters. Recall that each individual level variable influences the coefficient estimate through a multivariate regression. The individual-level variables are demeaned so that the regression intercepts reflect an "average" gambler.

INSERT TABLE 3 ABOUT HERE

The average gambler believes that that house advantage is around 52%. While higher than the true house advantage (about 12%) gamblers have substantial uncertainty surrounding this belief, with a standard deviation of .2. As expected, the cost coefficient is negative - high house advantage perceptions lower the probability of returning. The average gambler is risk seeking (at least directionally) and the offer values significantly influence the probability of returning. The interactions with the intercept are intuitive: gamblers that live far away are less likely to return while those in the higher tier LP are more likely to return. The posterior means for all of the parameters are presented in Table 4. The results for the fixed parameters are similar to the homogeneous results.

INSERT TABLE 4 ABOUT HERE

Table 5 displays the variances and correlations across the individual-level coefficient estimates. There is substantial heterogeneity across gamblers' coefficient estimates. Interestingly, there is a positive correlation between the prior mean and uncertainty: gamblers whose prior beliefs are higher tend to be more certain in their beliefs. There is also a strong negative correlation between the cost coefficient and the risk aversion; Gamblers who are more sensitive to the cost of gambling are more risk averse while those that are not as sensitive tend to be more risk seeking.

INSERT TABLE 5 ABOUT HERE

Figure 8 shows the distribution in posterior means across gamblers in the estimated prior house advantage and its uncertainty. Most players tend to overestimate the house advantage but the distribution is quite dispersed across gamblers. The level of uncertainty is somewhat bi-modal: while there is some mass around low uncertainty estimates there is also substantial mass around .05.

INSERT FIGURE 8 ABOUT HERE

Policy Simulations

The structural parameters are used to simulate six counterfactuals. The first two counterfactuals illustrate how projected casino revenues are quite sensitive to gamblers' prior beliefs in the house advantage and the volatility of outcomes. While these counterfactuals are informative, they do not provide casino marketers with practical solutions to act upon, for reasons to be discussed. The third and fourth counterfactuals focus on marketing solutions and show that sophisticated targeting strategies should consider how both the outcome sequence and prior beliefs may dictate where targeting is most effective. The remaining two counterfactuals explore belief-based targeting in more depth. The fifth counterfactual uses the model to identify the gamblers that are most responsive to marketing. Finally, the sixth counterfactual does a partial search for an optimal marketing strategy. While a full search is incredibly complex, the partial search still highlights that offer values should vary depending on both the outcome sequence and gambler beliefs.

Counterfactual 1: Accurate Prior Beliefs

In this data the average slot machine house advantage is 12.5%. The estimation results therefore suggest that gamblers overestimate the house advantage by a factor of about four prior to their first trip to the casino. Given that the cost coefficient θ_1 is negative, gamblers may be overestimating the cost of a return trip which in turn delays the return time. This counterfactual simulates expected gaming revenues under the assumption that each gambler's prior belief in the house advantage is accurate. That is, their prior belief equals the true house advantage. The results are shown in Table 6.

As expected, gamblers return at a faster rate if their prior beliefs in the house advantage are lower. With lower cost expectations gamblers no longer need many trips for their beliefs to converge to the true house advantage. Even though gamblers play less on each return trip the impact on the aggregate expected casino revenue is still positive.

If accurate beliefs in the house advantage can potentially increase long term casino revenue, why don't casino marketers simply advertise the accurate house advantages through direct mail? The primary reason is that this is not practical. Casinos tend to be very cautious on how they advertise slot machines in their direct mail offers. There is a risk that a gambler will interpret the true house advantage as a guaranteed loss limit. The casino may face backlash from the gamblers who lose more than the house advantage suggests they should. The purpose of presenting this counterfactual is to simply highlight that changes in a gambler's beliefs can have drastic long term consequences on casino revenues.

INSERT TABLE 6 ABOUT HERE

Counterfactual 2: Slot Machine Volatility

Next, I consider the impact of reducing the volatility of the slot machine hold variance. When a casino orders a slot machine from a manufacturer they specify the variability in that machine's outcomes. In this dataset, the slot machine hold is 13.9% and has a variance of .05, meaning 98% of the hold percentages (at the trip level) are between -38% and +66%. I simulate 1,000 gamblers over 5 years to measure the revenue impact of lowering and raising the hold variance relative to its current level. The results are presented in Table 7. Figure 9 plots the casino theoretical win against a multiplier on the hold variance - the dashed line at 1 means variance is at its current level.

The simulation results show that as the volatility decreases the projected casino win increases. However, when the volatility shrinks to a point that gambler wins become very infrequent the theoretical win declines. Clearly, the volatility in the outcomes has dramatic impacts on long term casino revenues. As with the first counterfactual, even though these findings are informative they do not point to any reasonable short term solution for managers. In order for a casino to change their aggregate slot machine volatility they would need to order new slot machines and spend time installing the machines on the gaming floor. These machine, labor, and additional opportunity costs are substantial and not accounted for here.

INSERT FIGURE 9 ABOUT HERE

INSERT TABLE 7 ABOUT HERE

Counterfactual 3: Incorporating Gambler Outcomes with Naive Targeting

The first two counterfactuals illustrate that changes in prior beliefs and hold percentage volatility can have substantial impacts on long term casino revenue. However, as discussed the results alone do not lend themselves immediately to practical solutions for managers. The purpose of these remaining four counterfactuals is to show how targeted marketing could be used in conjunction with the outcomes and player beliefs to improve casino profitability.

In this counterfactual I compare three marketing strategies: 1) the industry standard of basing offer values on gamblers' theoretical losses ("Industry Standard"), 2) basing offer values on actual outcomes but excluding gamblers players who won on their last trip ("Actual ex Wins"), and 3) basing offer values on theoretical losses (similar to the industry standard) but again excluding gamblers who won on their last trip ("Theo ex Wins"). The second and third strategies are meant to represent naive targeting strategies: gamblers who win are more likely to have low beliefs in the house advantage and therefore should be more likely to return to the casino anyways. Given this, the casino may be able to save on marketing expenses by excluding these players from offers. Furthermore, in the second scenario the casino provides an incentive to return that is directly in line with the loss experienced. I consider these strategies "naive" because they do not consider how each gambler's beliefs in the house advantage may influence the effectiveness of marketing - only the outcomes are used.

As in the empirical data, in each decision period there is a 45% chance that the gambler will be exposed to a marketing offer. The offers are valued at 30% of their last trip's theoretical or actual win. The total offer value is split into a room component and promotional credits, with two-thirds of the total offer value going to the room and one-third going to promotional credits. In the "Industry Standard" simulation, all gamblers have an opportunity to obtain an offer but in the "Actual ex Win" and "Theo ex Win" simulations offers will not be available to gamblers who won on their last trip.

The results in Table 8 show that both naive approaches to targeting are less profitable than the current industry standard. In the "Actual ex Wins" scenario, top line revenue remains relatively constant but the overall promotional costs are higher, even though fewer offers were redeemed. This may seem counterintuitive but this is because actual outcomes tend to have much more variability relative to theoretical outcomes, especially when evaluated at the trip level (as data becomes aggregated the theoretical outcomes converge to actual outcomes). In the "Theo ex Wins" strategy, promotional costs decrease dramatically but top line revenue also suffers. The short term gains that might be had from the reduction in promotional costs is offset by longer intervisit times. The results suggest that strategies that appear intuitive at first are not always more profitable in the long term. This counterfactual emphasizes the need for a more sophisticated targeting strategy.

INSERT TABLE 8 ABOUT HERE

Counterfactual 4: Marketing Impact by Past Outcomes and Beliefs

This simulation extends the previous by incorporating prior beliefs into the targeting decision. Table 9 shows the impact of marketing when gamblers' prior beliefs and uncertainty are high or low and when gamblers are either winning or losing. The impact of marketing is measured by comparing overall expected casino revenue with marketing versus without marketing. For example, an impact of .1 means that there is a 10% increase in revenue across gamblers in the presence of marketing. The marketing rule imposed is the same as the industry standard as described in the previous counterfactual. The "high" and "low" categorizations are set using the 5th and 95th percentile estimated prior means and prior variances.

When a gambler's prior belief in the house advantage is high and their uncertainty is high, marketing is more impactful if the player is on a *winning* streak rather than a losing streak. However, for gamblers whose prior beliefs in the house advantage are low, marketing is more impactful when players are on a *losing* streak. Notice that marketing is ineffective for gamblers whose beliefs in the house advantage are very high and their uncertainty is very low. This is intuitive: these gamblers are very certain that the cost is very high and because of they will not return regardless of marketing offers.

This simulation emphasizes the importance of considering both the prior beliefs *and* the outcome sequence when designing the targeting strategy. In the previous counterfactual the naive assumption was that only the outcome mattered but here we see that it is a combination of the outcomes and prior beliefs that dictate where marketing is more effective. This insight is very useful to managers who need to allocate their limited marketing budget across gamblers.

INSERT TABLE 9 ABOUT HERE

Counterfactual 5: Marketing Impact by Gambler

In this counterfactual I analyze the relationship between gamblers' *posterior* beliefs and the marketing impact. The posterior beliefs summarize both the prior beliefs and the outcome sequences realized, thereby reducing the number of metrics managers need to consider for targeting. To add more realism to this simulation, I use the 1,000 gamblers from the dataset rather than creating artificial gamblers. I simulate five years worth of gambling activity, picking up where the observed data ends. Again the focus is on the impact of marketing, meaning the change in expected casino revenue when there is marketing versus no marketing present. The goal of this simulation is to identify gamblers where marketing has the greatest impact and then determine if the marketing impact is in any way related to posterior mean and uncertainty in the belief of the house advantage.

Figure 10 shows the marketing impact represented by a lift chart. If gamblers were randomly targeted the total impact is expected to follow the dashed line. However, the simulations allow us to identify the gamblers where marketing will likely have the greatest impact.[8] Notice that just about all of the gains from marketing activity are realized from about one quarter of the gamblers. The other gamblers are not impacted by the marketing activity or in a few rare cases the marketing actions actually result in declines in gaming revenue.

INSERT FIGURE 10 ABOUT HERE

For gamblers who are most impacted by marketing (those in the front of the curve where the cumulative impact is less than 99%), the posterior belief in the house advantage tends to be higher and the uncertainty much lower.

	# of Gamblers	Posterior Mean	Posterior Uncertainty
99% of Cumulative Marketing Impact	242	.227	.0036
Remaining 1%	758	.196	.0076

Figure 11 illustrates the differences across gamblers. For each of the 1,000 gamblers, the marketing impact is plotted against the posterior mean and posterior uncertainty averaged across all of their realized return trips. Marketing has a greater impact on gamblers with higher posterior means and lower uncertainty. The correlation between the posterior mean and posterior variance across gamblers is -.128: gamblers with higher beliefs in the house advantage tend to have less uncertainty. While this may seem to contradict the findings from the previous counterfactual it is important to note that the previous counterfactual examined the extremes of beliefs, at the 5th and 95th percentile. In addition, this counterfactual uses the actual gamblers, rather than simulated gamblers. In both cases the fact remains that there is a strong relationship between the posterior beliefs and the impact of marketing.

INSERT FIGURE 11 ABOUT HERE

Counterfactual 6: Optimal Marketing Offers

The previous counterfactual provides evidence that posterior beliefs influence the impact of marketing. A natural extension is to then search for the optimal marketing strategy. That is, for each gambler and each

[8] Since the impact of marketing depends on the outcome sequence a more thorough analysis would simulate over many potential outcome paths. I conducted a simulation setting the hold percentage to a constant (the mean) and the interpretations are the same.

outcome experience which offer strategy will lead to the highest long term expected revenue? Finding the global optimum is very difficult (at least in this casino example) because each room value and slot promotional credit combination would need to be evaluated for each gambler at each decision period for every potential outcome sequence. Even though finding the global optimum is incredibly complex this counterfactual shows that even a relatively simple constrained optimization can lead to substantial improvements in projected revenue.

In this constrained search I vary the slot promotional credits and bin the posterior beliefs into four categories. The goal is to determine how much each of the four posterior belief categories should receive in slot promotional credits. In this dataset the promotional credit value is typically set at 10% of the past theoretical loss level. I simulate this baseline percentage and four alternatives: 0%, 5%, 15%, and 20%. The belief and uncertainty levels are grouped into four categories: high/low belief in the house advantage and high/low uncertainty. The cutoff for the belief in the house advantage is the casino's true house advantage and the cutoff for the uncertainty is based on a median split of the observed gambler's posterior variances.

Category	Belief in House Advantage	Uncertainty in Belief
Low	<12.5%	<.0029
High	>=12.5%	>=.0029

Another challenge in searching for the optimal marketing offer is that gamblers can switch categories over time depending on their outcomes. That is, they may start in a high belief/high uncertainty state, move to high belief/low uncertainty state, and then end in a low belief/low uncertainty state. Because each state will have its own marketing strategy, all 625 combinations of offers need to be considered: five promotional credit percentages in each of the four offer states.

I simulate one hundred gamblers for two years in each of the 625 offer value combinations. Each gambler starts with the same prior beliefs and uncertainty on the house advantage (based on the hierarchical results for the "average" gambler). Profit is obtained by subtracting room and promotional credit costs from the projected casino revenue.

Figure 12 shows the sorted profit across all 625 simulations. The dashed line shows baseline profitability where the four belief categories each receive promotional credits valued at 10% of theoretical losses. The range in the profit is substantial: the top strategies generate over $55,000 in profit while the worst strategies generate around $25,000.

INSERT FIGURE 12 ABOUT HERE

Rather than try to evaluate each of the 625 simulations individually, I instead compare the differences in the most and least profitable strategies, shown in Figure 13. This shows which promotional credit percentage is associated with the most and least profitable strategies in each of the four belief/uncertainty categories. Notice that the most profitable strategy does not use the baseline percentage of 10% in any of the four belief categories: when the belief in the house advantage is below the actual house advantage, a higher percentage is recommended whereas when the belief in the house advantage is high the policy depends on the uncertainty. It is also interesting to note that the most profitable strategy does not max out or eliminate the promotional credit amount in any of the four belief bins, suggesting that the solution is contained within the boundaries of the simulation. The most profitable strategy generated $62,290 in profit, compared to $36,479 in the baseline scenario where all gamblers receive the same promotional credit percentage regardless of their beliefs in the house advantage, an increase of 85.3%. For a more conservative (and realistic) measure of success, the top half of the strategies still increased baseline profit by an average of 19.7%.

The model presented provides a framework for managers to use in order to target gamblers based on their beliefs and outcome sequences. This simulation shows that the gains from doing so can be significant, even when the strategy employed is the result of a heavily constrained search.

INSERT FIGURE 13 ABOUT HERE

Discussion and Conclusion

This paper develops a dynamic Bayesian learning model for the purpose of assisting managers with direct marketing decisions when the customer's experience quality is observed by the firm. This class of problem arises naturally in many industries, mostly those that are service-based where experience variability is often unavoidable across occasions. Depending on the signal variability, it may take many experiences for the customer to learn the true distribution of quality. Until the true distribution is learned, potentially biased perceptions may warrant additional targeted marketing, especially if the quality of a customer's initial signals are atypically low. To facilitate tractable estimation I take advantage of inexpensive cloud computing and exploit the massively parallel structure of combining forward simulation with a utility function that is linear in parameters. The proposed structure easily incorporates flexible heterogeneity distributions to generate individual level parameter estimates, which is central to many modern targeted marketing problems.

The model is estimated on data from a casino where gamblers learn about the average slot machine house advantage. Gamblers use their beliefs on the house advantage to project future trip costs which in turn influences when they return to the casino and how they play on a return trip. The gaming industry is an ideal setting to study this model for a variety of reasons but primarily because the exogenous gambling outcomes provide many distinct and unique experience sequences at the gambler level. The results and counterfactuals suggest that gamblers tend to overestimate the house advantage, which increases projections in gambling expenses and delays in the time until the next trip. The counterfactuals also highlight the importance of incorporating gambler beliefs and outcomes into the marketing decisions. The simulations suggest that CRM managers can benefit tremendously when customer beliefs and outcome sequences are considered in designing one-to-one targeting strategies.

There are a few limitations worth addressing in regards to the estimation strategy and the empirical analysis. While the estimation strategy does improve the tractability substantially, given that it is not a full solution approach it is not immediately clear the number of paths to forward simulate for each starting state to ensure the state space is sufficiently explored. In my estimation I continued to simulate additional paths until the discounted future values appeared to converge - I then doubled this number of simulations as an extra precaution in the final estimation. Another limitation previously mentioned is that the discount factor cannot be estimated, but again it is relatively easy to test multiple discount factors by taking advantage of the parallel design. One limitation with the empirical analysis is that competitor activity is not observed. This includes both player activity at competitor casinos and competitor marketing activity. For instance, I do not know if delays in return trips are due to gamblers visiting other casinos, the marketing activities of competitor casinos, or simply a lack of gambling. A more central concern is that gamblers may be learning from slot machines at outside casinos between the casino trips I observe. At first glance, multi-casino gambling appears to impact the learning process substantially. However, it is important to note that each casino will have its own mix of slot machines, which means the average house advantage of each casino is likely to be different. Even if a gambler visits other casinos, they still need to learn about *this* casino's house advantage. Finally, there are characteristics on the counterfactuals that warrant discussion. First, I don't account for competitor reactions. However, competitor reactions are unlikely to be much of a concern because 1) competitors will not know which gamblers have been identified as being responsive to marketing, and 2) competitors don't know the outcome sequences experienced by gamblers at this casino. In other words, given the targeting strategies presented it is not immediately clear how a competitor could react. Second, I don't allow gamblers to learn about the targeting strategy over time. However, since the relatively simple targeting strategies were shown to be ineffective a gambler would need to learn a very sophisticated strategy, which is quite unlikely unless they have an unrealistically high number of trips and marketing exposures.

There are many possible extensions to this current work. One important extension is to allow a learning rate to be estimated at the individual level. In this paper the learning rate is fixed due to the formulation of the Bayesian updating process but by including this additional parameter gamblers can potentially learn at faster or slower rates, similar to Narayanan and Manchanda (2009) but with forward looking consumers. The speed of learning could have substantial impacts on the targeted marketing decisions. Another extension could account for how beliefs influence projected marketing offers which in turn may change current period decisions. Finally, there may be more efficient ways to search for the global optimal marketing strategy when using this model for targeting. I leave these topics for future research.

References

American Gaming Association (2013). 2013 State of the States: The AGA Survey of Casino Entertainment.

Bajari, P., Benkard, C. L., and Levin, J. (2007). Estimating dynamic models of imperfect competition. *Econometrica*, 75(5):1331–1370.

Bolton, R. N. (1998). A dynamic model of the duration of the customer's relationship with a continuous service provider: The role of satisfaction. *Marketing science*, 17(1):45–65.

Boulding, W., Kalra, A., Staelin, R., and Zeithaml, V. A. (1993). A dynamic process model of service quality: from expectations to behavioral intentions. *Journal of marketing research*, 30(1):7.

Chib, S. and Greenberg, E. (1996). Markov chain monte carlo simulation methods in econometrics. *Econometric theory*, 12(03):409–431.

Ching, A. T., Erdem, T., and Keane, M. P. (2013). Invited paper-learning models: An assessment of progress, challenges, and new developments. *Marketing Science*, 32(6):913–938.

Compton, J. (1999). Players clubs 101.

Eckstein, Z., Horsky, D., and Raban, Y. (1988). An empirical dynamic model of optimal brand choice. *Foerder Institute of Economic Research, Working Paper*, (88).

Erdem, T. and Keane, M. P. (1996). Decision-making under uncertainty: Capturing dynamic brand choice processes in turbulent consumer goods markets. *Marketing Science*, 15(1):1–20.

Hotz, V. J., Miller, R. A., Sanders, S., and Smith, J. (1994). A simulation estimator for dynamic models of discrete choice. *The Review of Economic Studies*, 61(2):265–289.

Imai, S., Jain, N., and Ching, A. (2009). Bayesian estimation of dynamic discrete choice models. *Econometrica*, 77(6):1865–1899.

Keane, M. P. and Wolpin, K. I. (1994). The solution and estimation of discrete choice dynamic programming models by simulation and interpolation: Monte carlo evidence. *The Review of Economics and Statistics*, 76(4):648–672.

Kilby, J. and Fox, J. (1998). *Casino Operations*. Wiley & Sons.

Li, X. (2014). Product offerings and product line length dynamics. *Available at SSRN 2462577*.

Meyer, R. J. (1981). A model of multiattribute judgments under attribute uncertainty and informational constraint. *Journal of Marketing Research*, pages 428–441.

Nair, H., Misra, S., Hornbuckle, W. J., Mishra, R., and Acharya, A. (2013). Big data and marketing analytics in gaming: Combining empirical models and field experimentation.

Narayanan, S. and Manchanda, P. (2009). Heterogeneous learning and the targeting of marketing communication for new products. *Marketing Science*, 28(3):424–441.

Narayanan, S. and Manchanda, P. (2012). An empirical analysis of individual level casino gambling behavior. *Quantitative Marketing and Economics*, 10(1):27–62.

Osborne, M. (2011). Consumer learning, switching costs, and heterogeneity: A structural examination. *Quantitative Marketing and Economics*, 9(1):25–70.

Oxford Economics (2014). Economic impact of the us gaming industry.

Park, H. M. and Manchanda, P. (2015). When harry bet with sally: An empirical analysis of multiple peer effects in casino gambling behavior. *Marketing Science*, 34(2):179–194.

Reiss, P. C. (2011). Structural workshop paper-descriptive, structural, and experimental empirical methods in marketing research. *Marketing Science*, 30(6):950–964.

Roberts, J. H. and Urban, G. L. (1988). Modeling multiattribute utility, risk, and belief dynamics for new consumer durable brand choice. *Management Science*, 34(2):167–185.

Roos, J. M., Mela, C. F., and Shachar, R. (2013). Hyper-media search and consumption. *Available at SSRN 2286000*.

Rossi, P. (2014). *Bayesian Non-and Semi-parametric Methods and Applications*. Princeton University Press.

Rossi, P. E., Allenby, G. M., and McCulloch, R. E. (2005). *Bayesian Statistics and Marketing*. John Wiley & Sons.

Rust, R. T., Inman, J. J., Jia, J., and Zahorik, A. (1999). What you don't know about customer-perceived quality: The role of customer expectation distributions. *Marketing Science*, 18(1):77–92.

Sriram, S., Chintagunta, P. K., and Manchanda, P. (2015). Service quality variability and termination behavior. *Management Science*, 61(11):2739–2759.

Tversky, A. and Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *science*, 185(4157):1124–1131.

Van den Berg, G. J. (2001). Duration models: specification, identification and multiple durations. *Handbook of econometrics*, 5:3381–3460.

Tables and Figures

Table 1: DESCRIPTIVE STATISTICS

	All Data	Gamblers >= 3 Trips
Gamblers	28,362	13,964
Trips	113,752	94,139
Trip Length	2.5 days	2.8 days
Slot Average Bet	$1.95	$1.69
Spins/Minute	7.7	7.8
Hours Played/Trip	5.3	5.9
Hold % Experienced	12%	10%
Expected Loss/Trip	$387	$416

Table 2: HOMOGENEOUS RESULTS

Coefficient	Posterior Mean	SE	Coefficient	Posterior Mean	SE
Intercept	-2.8563	(1.69e-03)	γ_1	-0.3543	(1.88e-03)
A_0	0.3973	(1.54e-04)	γ_2	-0.1600	(1.82e-03)
σ_0^2	0.0012	(2.89e-06)	γ_3	-0.1878	(1.73e-03)
Cost	-0.0018	(1.72e-06)	γ_4	-0.6775	(1.88e-03)
Risk	-1.85e-07	(1.12e-09)	γ_5	-0.0153	(1.94e-03)
Gaming Offer	0.0050	(5.46e-06)	γ_6	0.0484	(1.77e-03)
Room Offer	0.0168	(5.03e-06)	γ_7	0.0775	(1.77e-03)
ω_1	0.0372	(2.54e-05)	γ_8	0.2299	(1.78e-03)
ω_2	-0.0013	(2.20e-07)	γ_9	0.2059	(1.76e-03)
ω_3	1.31e-05	(5.78e-10)	γ_{10}	-0.1051	(1.80e-03)
ω_4	-5.10e-08	(3.25e-15)	γ_{11}	-0.0606	(1.78e-03)
ω_5	6.78e-11	(2.68e-17)			

Table 3: HIERARCHICAL INTERACTIONS

		Demographics				
Coefficient	Description	Intercept	Age (divided by 10)	Male	Log Distance (miles)	Gold LP Card
A_0	Prior mean	0.523*	-0.014	-0.004	-0.002	-0.003
σ_0^2	Prior uncertainty	0.040*	0.000	0.001	0.001	-0.005
θ_1	Cost	-0.067*	-0.011	0.012	-0.008	-0.035
r	Risk	-0.004	0.001	-0.002	0.002	-0.002
θ_2	Offer promo credits	0.151*	0.024	0.004	0.012	0.118
θ_3	Offer room value	0.075*	-0.048	0.011	-0.141*	0.008
θ_0	Intercept	-2.793*	0.064	0.189*	-0.092*	0.287*

* = 95% highest posterior density does not cover zero

Table 4: FULL HIERARCHICAL RESULTS

Coefficient	Description	Posterior Mean	SE
A_0	Prior mean	0.5231	(2.40e-04)
σ_0^2	Prior uncertainty	0.0398	(2.30e-05)
θ_1	Cost	-0.0668	(1.24e-04)
r	Risk	-3.76e-03	(7.84e-05)
θ_2	Offer promo credits	0.1510	(3.88e-04)
θ_3	Offer room value	0.0765	(3.05e-04)
θ_0	Intercept	-2.7933	(4.98e-04)
ω_1	Weeks since last trip^1	0.0489	(1.82e-05)
ω_2	Weeks since last trip^2	-0.0013	(1.57e-07)
ω_3	Weeks since last trip^3	1.31e-05	(1.22e-10)
ω_4	Weeks since last trip^4	-5.10e-08	(1.59e-15)
ω_5	Weeks since last trip^5	6.78e-11	(1.22e-17)
γ_1	Jan	-0.6550	(1.90e-03)
γ_2	Feb	-0.4129	(1.72e-03)
γ_3	Mar	-0.5658	(1.75e-03)
γ_4	Apr	-0.7259	(1.85e-03)
γ_5	May	-0.2667	(1.75e-03)
γ_6	Jun	-0.1619	(1.58e-03)
γ_7	Jul	-0.1741	(1.76e-03)
γ_8	Aug	-0.0560	(1.63e-03)
γ_9	Sep	-0.0990	(1.70e-03)
γ_{10}	Oct	-0.2499	(1.77e-03)
γ_{11}	Nov	-0.1716	(1.63e-03)

Table 5: HETEROGENEITY ACROSS GAMBLERS

Coefficient	Description	Variance (diagonal) and Correlation (off-diagonal)						
A_0	Prior mean	.1337						
σ_0^2	Prior uncertainty	-.24	.0015					
θ_1	Cost	.06	-.03	.0563				
r	Risk	-.01	.02	-.19	.0142			
θ_2	Offer promo credits	.00	-.06	-.05	-.04	.7150		
θ_3	Offer room value	-.01	.05	.02	.05	-.55	.8034	
θ_0	Intercept	.13	-.10	.00	-.02	.06	-.06	1.0623

Table 6: ACCURATE PRIORS INCREASE CASINO REVENUE

	Current Prior	Accurate Prior
A_0	.523	.125
Trips	2,379	10,027
Average Weeks to Next Trip	24	16
Average Trip Slot Theoretical Loss	$460	$162
Total Theoretical Loss	$1,094,114	$1,621,907
Increase in Gaming Revenue	48.2%	
# of Gamblers Simulated	1,000	
Years Simulated	5	

Table 7: HOLD PERCENTAGE VOLATILITY IMPACTS CASINO REVENUES

Variance Multiplier	Variance	1% LB	99% UB	Avg. Return Weeks	Avg. Theo. Loss	Total Theo. Loss
0.001	.0002	.09	.16	29	$362	$2,630,910
0.005	.0010	.05	.20	26	$353	$2,861,178
0.010	.0021	.02	.23	23	$345	$3,138,954
0.025	0.01	-.04	.30	16	$323	$4,193,657
0.050	0.01	-.11	.37	11	$299	$5,925,798
0.200	0.04	-.35	.60	10	$319	$5,403,860
0.350	0.07	-.50	.76	15	$374	$2,672,789
0.500	0.10	-.62	.88	19	$410	$1,626,595
0.650	0.14	-.73	.99	23	$442	$1,145,819
0.800	0.17	-.82	1.08	24	$446	$913,859
0.950	0.2	-.91	1.17	23	$459	$761,572
1.100	0.23	-.99	1.24	28	$476	$621,727
1.250	0.26	-1.06	1.32	25	$465	$563,224
1.400	0.29	-1.13	1.39	26	$481	$556,337

Table 8: NAIVE TARGETING IS INEFFECTIVE

Targeting Criteria	Industry Standard	Actual ex Wins	Theo. ex Wins
Trips	3,084	2,441	2,388
Avg. Weeks to Return	19	25	24
Avg. Theoretical Win/Trip	$472	$498	$482
Total Theoretical Win	$1,456,769	$1,215,938	$1,150,119
Total Actual Win	$1,453,383	$1,408,149	$1,266,449
Promotions Redeemed	1,328	724	697
Room Value	$128,509	$168,937	$65,902
Promotional Credits	$64,254	$84,469	$32,951
Room Cost ($30 per roomnight)	$43,380	$41,460	$21,840
Promotional Credit Cost (1 cycle)*	$56,222	$73,910	$28,832
Net Theoretical Win	$1,357,167	$1,100,568	$1,099,446
Net Actual Win	$1,353,780	$1,292,779	$1,215,777

*The cost of promotional credits is not a certainty since wins can be cycled back into the machine and generate additional payouts. See the Web Appendix for a discussion.

Table 9: MARKETING IMPACT DEPENDS ON PRIOR BELIEFS

Prior Belief in House Advantage	Prior Uncertainty	Player Winning	Player Losing	Δ
High	High	3.8	1.3	2.6
High	Low	0.0	0.0	0.0
Low	High	0.1	3.4	-3.3
Low	Low	0.3	2.4	-2.1

Figure 1: INTERVISIT TIMING DISTRIBUTION

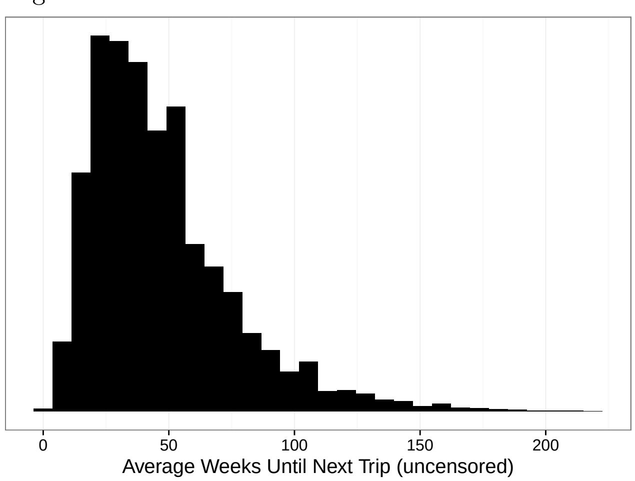

Figure 2: MEDIAN RETURN LAG AFTER A PLAYER LOSS (RELATIVE TO A PLAYER WIN)

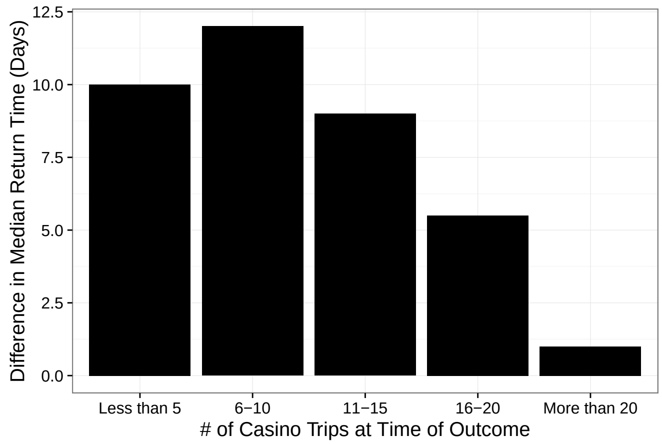

Figure 3: RETURN LAG BY PLAYER LOSING STREAK (RELATIVE TO PLAYER WINNING STREAK)

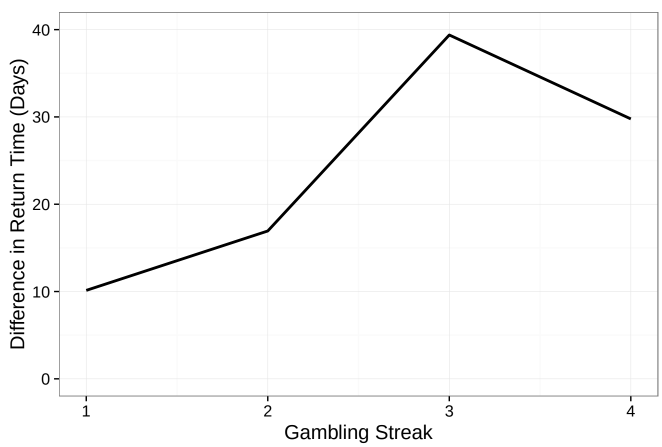

Figure 4: TRUNCATED NORMAL DISTRIBUTION EXAMPLES USING GRIDPOINTS

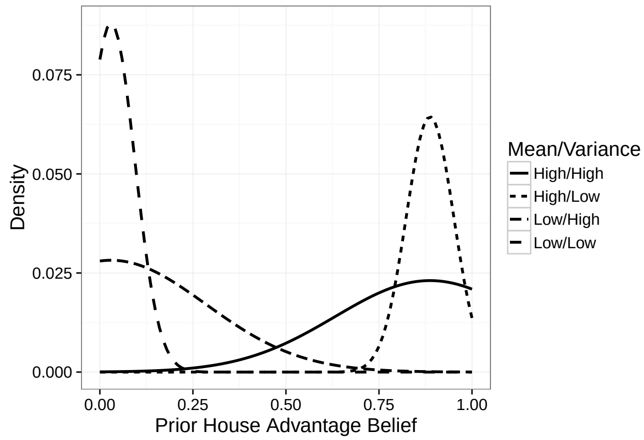

Figure 5: WEIBULL POSTERIOR MEAN COEFFICIENTS

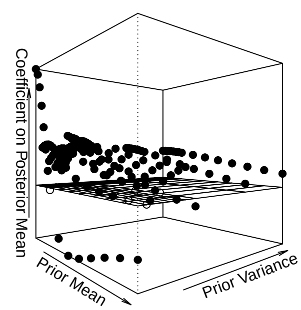

Figure 6: WEIBULL POSTERIOR VARIANCE COEFFICIENTS

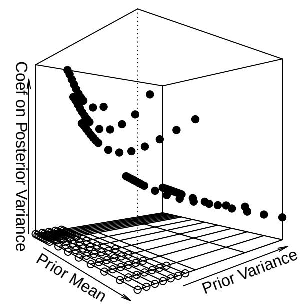

Figure 7: GAMBLERS LEARN THE HOUSE ADVANTAGE FROM THE HOLD PERCENTAGES EXPERIENCED

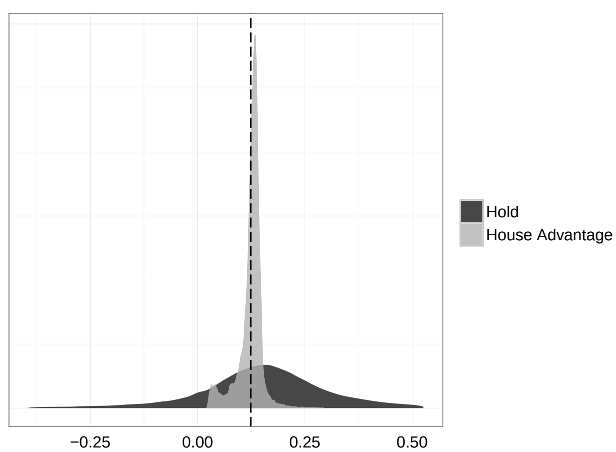

Figure 8: HETEROGENEITY IN PRIOR BELIEFS AND UNCERTAINTY ACROSS GAMBLERS

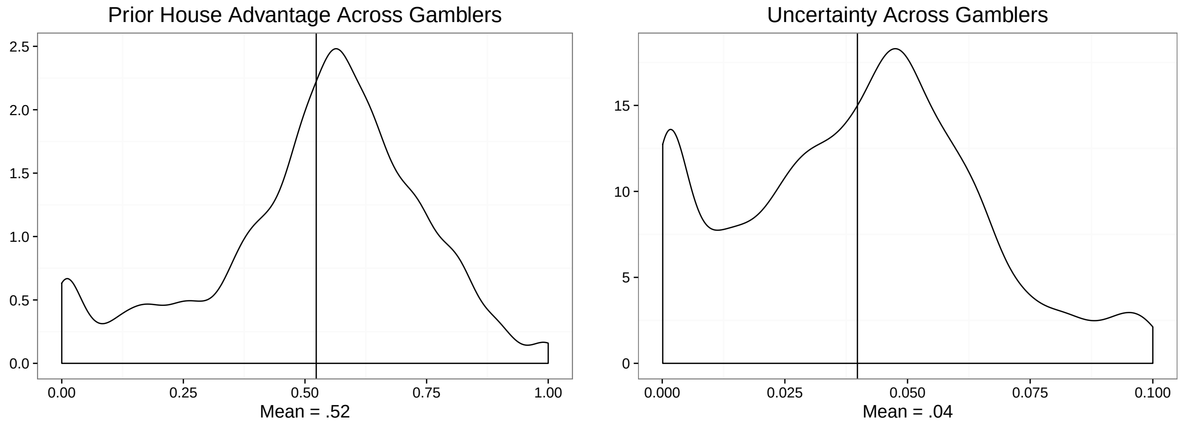

Figure 9: IMPACT OF HOLD PERCENTAGE VOLATILITY ON EXPECTED CASINO REVENUE

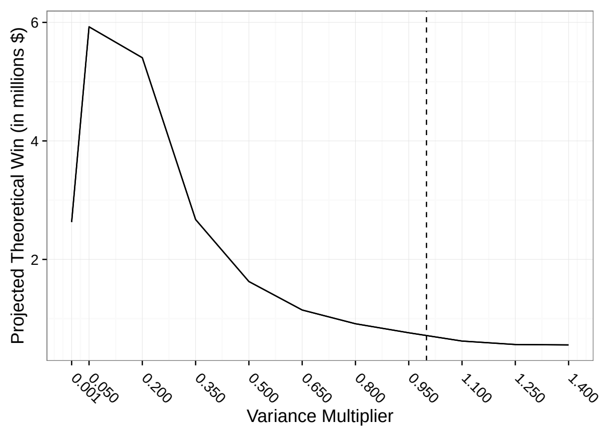

Figure 10: TARGETED MARKETING LIFT

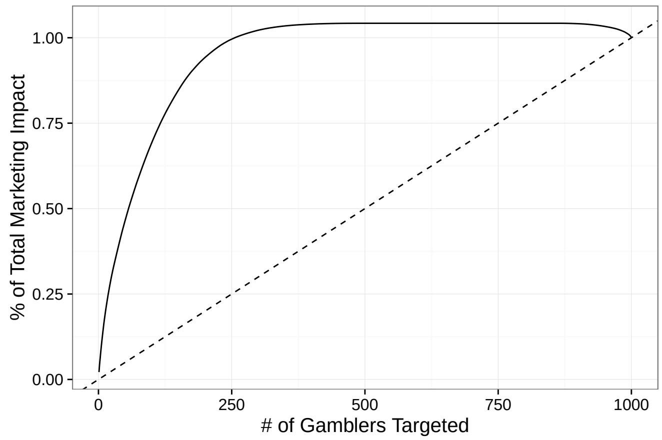

Figure 11: MARKETING IMPACT BY POSTERIOR BELIEFS AND UNCERTAINTY ACROSS GAMBLERS

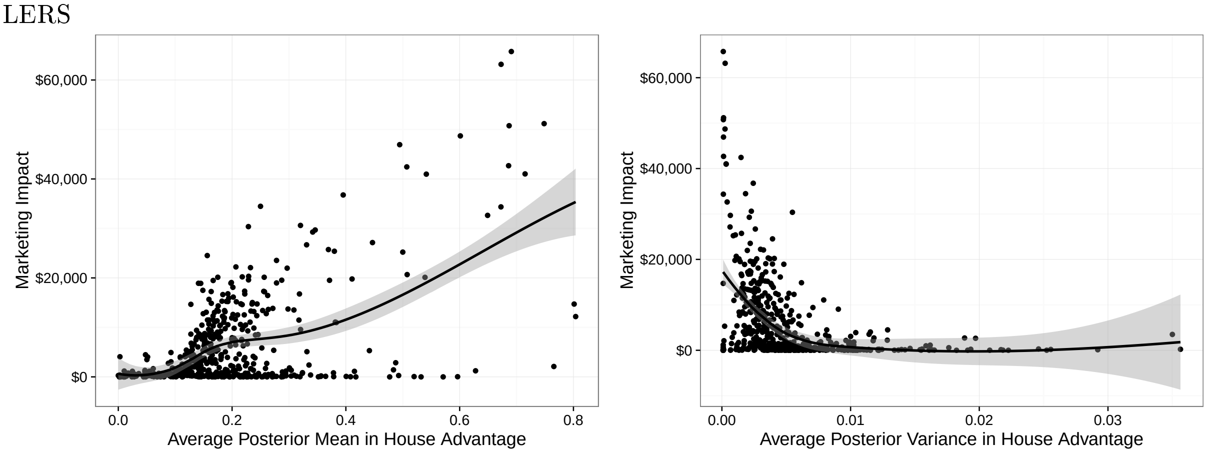

Figure 12: SIMULATED CASINO PROFIT BY TARGETING STRATEGY

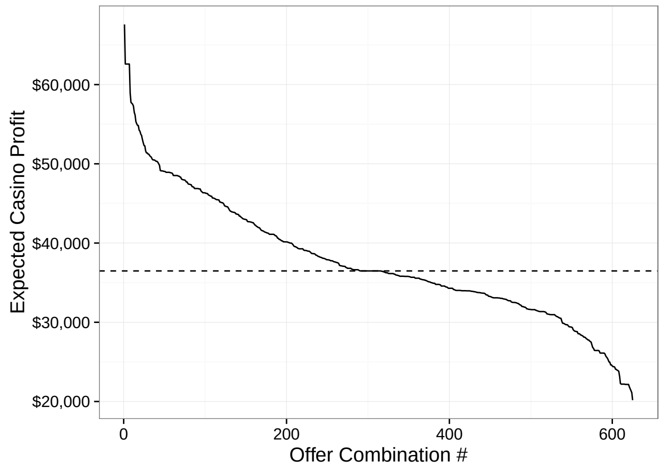

Figure 13: CONSTRAINED OPTIMIZATION BEST AND WORST STRATEGY PROFILES

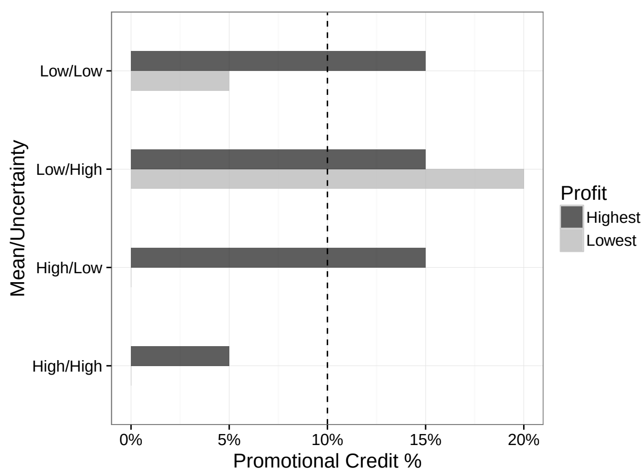

Appendix A. Conjugate Prior in Truncated Normal Distribution

If the prior is truncated normal and the signal is an unbounded normal then the corresponding posterior is also a truncated normal. In other words, the truncated normal distribution is also a conjugated prior for a standard normal likelihood of signal generation. This proof is similar to the one in Li (2014).

Theorem 1. Suppose the parameter of interest θ is distributed in normal distribution truncated at 0 and 1, i.e., $\theta \sim \mathcal{TN}\left(\mu_0, \sigma_0^2 = \lambda_0^{-1}, 0, 1\right)$, and the likelihood for signal

$$x = \theta + \xi$$

where $\xi \sim \mathcal{N}\left(0, \sigma_\xi^2 = \lambda_\xi^{-1}\right)$, then the posterior distribution

$$\theta | x \sim \mathcal{TN}\left(\mu_1, \sigma_1^2 = \lambda_1^{-1}, 0, 1\right)$$

with

$$\begin{aligned} \mu_1 &= \frac{\lambda_0}{\lambda_0 + \lambda_\xi}\mu_0 + \frac{\lambda_\xi}{\lambda_0 + \lambda_\xi}x \\ \lambda_1 &= \lambda_0 + \lambda_\xi \end{aligned}$$

Proof. Let $\phi\left(t, \mu, \sigma^2\right)$ be the normal pdf with mean μ and variance σ^2, and $\Phi\left(t, \mu, \sigma^2\right) = \int_{-\infty}^{t} \phi\left(s, \mu, \sigma^2\right) ds$ be the CDF. We know that

$$\begin{aligned} f(\theta) &= \frac{\phi\left(\theta, \mu_0, \sigma_0^2\right)}{\Phi\left(1, \mu_0, \sigma_0^2\right) - \Phi\left(0, \mu_0, \sigma_0^2\right)} \\ f(x|\theta) &= \phi\left(x, \theta, \sigma_\xi^2\right) \end{aligned}$$

so

$$f(x) = \int_0^1 f(x|\theta) f(\theta) d\theta = \frac{\int_0^1 \phi\left(x, \theta, \sigma_\xi^2\right) \phi\left(\theta, \mu_0, \sigma_0^2\right) d\theta}{\Phi\left(1, \mu_0, \sigma_0^2\right) - \Phi\left(0, \mu_0, \sigma_0^2\right)}$$

and

$$
\begin{aligned}
f(\theta|x) &= \frac{f(\theta) f(x|\theta)}{f(x)} = \frac{\phi\left(x,\theta,\sigma_\xi^2\right)\phi\left(\theta,\mu_0,\sigma_0^2\right)}{\int_0^1 \phi\left(x,\theta,\sigma_\xi^2\right)\phi\left(\theta,\mu_0,\sigma_0^2\right)d\theta} \\
&= \frac{\phi\left(x,\theta,\sigma_\xi^2\right)\phi\left(\theta,\mu_0,\sigma_0^2\right)/\int_{-\infty}^{\infty}\phi\left(x,\theta,\sigma_\xi^2\right)\phi\left(\theta,\mu_0,\sigma_0^2\right)d\theta}{\int_0^1 \phi\left(x,\theta,\sigma_\xi^2\right)\phi\left(\theta,\mu_0,\sigma_0^2\right)d\theta/\int_{-\infty}^{\infty}\phi\left(x,\theta,\sigma_\xi^2\right)\phi\left(\theta,\mu_0,\sigma_0^2\right)d\theta} \\
&= \frac{\phi\left(\theta,\mu_1,\sigma_1^2\right)}{\int_0^1 \phi\left(\theta,\mu_1,\sigma_1^2\right)d\theta} = \frac{\phi\left(\theta,\mu_1,\sigma_1^2\right)}{\Phi\left(1,\mu_1,\sigma_1^2\right)-\Phi\left(0,\mu_1,\sigma_1^2\right)}
\end{aligned}
$$

where the second equality before last is obtained by the standard conjugate prior of the normal distribution, which is

$$
\begin{aligned}
\theta &\sim \mathcal{N}\left(\mu_0,\sigma_0^2\right) \\
x|\theta &\sim \mathcal{N}\left(\theta,\sigma_\xi^2\right)
\end{aligned}
$$

will imply

$$
\theta|x \sim \mathcal{N}\left(\mu_1,\sigma_1^2\right)
$$

Appendix B. Forward Simulation Algorithm

```
Algorithm: Forward Simulation Pseudo-Code

for simulation r ← 1 to R do
    foreach starting state s_0 do
        foreach starting action a_0 ∈ {0, 1} do
            for time period t ← 0 to ∞ do
                if t_0 = 0 then
                    returnFlag = a_0
                else if t_0 > 0 then
                    Calculate probability of returning based on state
                    Draw from uniform [0,1] and update returnFlag based on probability
                end
                Record the current state values (to be discounted)
                if returnFlag = 1 then
                    Reset the weeks since last trip
                    Increment the month (if needed)
                    Draw a new hold value
                    Update beliefs based on hold
                else if returnFlag = 0 then
                    Increment week counter
                    Increment the month (if needed)
                end
                Update play style for next trip
                Draw a new marketing offer
            end
            Discount the values by β
        end
    end
end
```

Algorithm 1: Forward Simulation