WEB APPENDIX A. COST OF PROMOTIONAL CREDITS

The casino industry remains divided on the cost of promotional credits. The reason is because the true cost depends on how many times the credit is cycled back through the slot machine, as illustrated below. In this example, a \$100 promotional credit is played through a slot machine with a 12.5% house advantage. That means the casino is expected to keep \$12.50 and pay out \$87.50, which the gambler can keep and convert to cash. However, if the gambler cycles their winnings back into the machine the casino is expected to keep another 12.5% (or \$10.94) and pay back \$76.56, which is now the true cost of the promotional credit. Obviously, this can continue until the gambler has cycled the all of the winnings through the machine, resulting in zero cost to the casino. It becomes even more difficult because not all slot machines have the same house advantage, and winnings from one machine can be played on another. Because of these complexities, financial planning departments often use simple rules such as charging 100% of the cost or some fixed fraction of projected promotional credits redeemed.

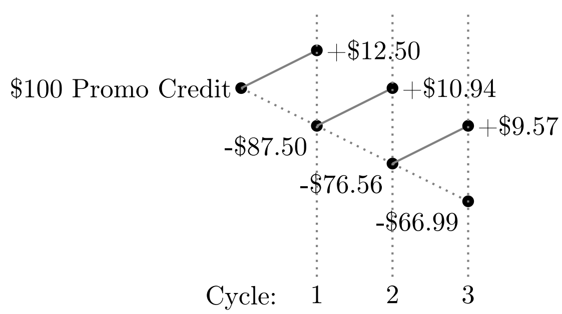

WEB APPENDIX B. BI-LINEAR INTERPOLATION

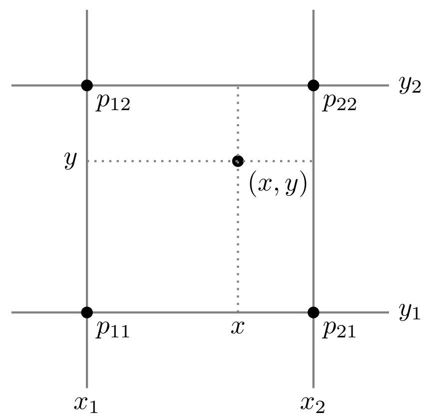

Bi-linear interpolation is an extension of linear interpolation. Linear interpolation is performed in one direction, and then again in the other. In this analysis, the learning process prior mean and variance gridpoints form the bi-linear lattice. The forward simulation yields discounted expected values at each prior mean and prior variance gridpoint (labeled p_{11}, p_{12}, p_{21}, and p_{22} in the figure above). During the MCMC draws, a new prior mean and prior variance may be suggested that has not been forward simulated over, in this example point (x, y). The below formula is used to approximate that point's discounted expected value, $f(x, y)$.

$$f(x,y) = \frac{f(p_{11})(x_2 - x)(y_2 - y) + f(p_{21})(x - x_1)(y_2 - y) + f(p_{12})(x_2 - x)(y - y_1) + f(p_{22})(x - x_1)(y - y_1)}{(x_2 - x_1)(y_2 - y_1)}$$

WEB APPENDIX C. ESTIMATION PROCEDURE

Step 0. Forward Simulation Settings

The casino data is constructed so that each record represents a decision period. The first record for each gambler is the first trip to the casino and each week after the first trip is a new record. In preparation for the forward simulation, the discount factor, number of periods to forward simulate, and number of forward simulations per starting state need to be selected.

Parameter	Description	Value
β	discount factor	.98
T	time periods to forward simulate	228
R	number of forward simulations per starting state	100

I selected T such that it is the smallest value for which β^T is less than .01 so that simulating beyond this point will have negligible impacts on the discounted values. I selected R based on experimentation; I found that at around $R = 50$ the discounted values began to converge. As a conservative measure I then doubled this to $R = 100$.

The forward simulation is done across 150 prior mean and prior variance combinations. This consists of 15 prior mean values and 10 prior variance values, given below. As illustrated in the paper, the gridpoints were selected in order to represent a broad range of truncated normal distributions.

A_0 Grid	σ_0^2 Grid
0.0000010	0.0001
0.0300010	0.0041
0.0600010	0.0081
0.0900010	0.0121
0.1200010	0.0161
0.1500010	0.0201
0.1800010	0.0241
0.2100000	0.0300
0.3228571	0.0650
0.4357143	0.1000
0.5485714	
0.6614286	
0.7742857	
0.8871429	
1.0000000	

Step 1. Play Style and Policy Function Estimation

For each gridpoint, I estimate the coefficients for the play style regression and the non-parametric policy function. These estimates need to be done for each gridpoint because each prior mean and prior variance changes how the posterior beliefs evolve. In other words, the coefficient estimates depend on the learning prior settings. The play style priors are set to the following:

$$
\begin{aligned}
\bar{\beta} &= 0 \\
A &= .01 \\
\nu &= 3 \\
\text{ssq} &= \text{var}(y)
\end{aligned}
$$

Here is a summary the posterior means of the play style coefficient estimates across the 150 gridpoints:

	β_1	β_2	β_3	β_4	β_5	β_6	β_7	β_0
25%	1.5203	-3.3670	-0.00138	0.00151	-0.00145	0.00163	-0.00048	-0.25878
50%	2.2112	2.0343	-0.00135	0.00154	-0.00140	0.00164	-0.00048	-0.19895
75%	2.5324	13.4973	-0.00133	0.00155	-0.00137	0.00165	-0.00046	-0.14189
Min	0.0706	-8.5708	-0.00155	0.00139	-0.00159	0.00160	-0.00049	-0.36145
Median	2.2112	2.0343	-0.00135	0.00154	-0.00140	0.00164	-0.00048	-0.19895
Mean	2.0307	6.6500	-0.00137	0.00153	-0.00142	0.00164	-0.00047	-0.18996
Max	3.7892	51.9811	-0.00130	0.00157	-0.00136	0.00167	-0.00043	0.17292

Details of the non-parametric policy function estimation are given in the paper. The prior settings used are as follows:

$$\begin{aligned} \alpha &= 5 \\ \bar{\mu} &= 0 \\ a_u &= .01 \\ \nu &= 8 \\ V &= \operatorname{diag}(\nu) \end{aligned}$$

Step 2. Forward Simulation

At this point the data is ready for forward simulation. I divided the data into 200 groups, each group contained the data associated with 5 gamblers. I sent each of the 200 groups with each of the 150 gridpoints to Amazon EC2, meaning a total of 30,000 files would be forward simulated over and then collected. The parallelization can be as granular as one wishes however since Amazon rounds up on each instance hour consumed I divided the data based on how many paths I could comfortably process in one hour (allowing for time to send and receive the data from Amazon via ssh). Otherwise large amounts of processing time (and cost) would be wasted. The sending and receiving of data to Amazon was done automatically through a shell script - the instances were stagger started to avoid overloading the ssh calls and data transfers.

For a single simulation (i.e., 1 of the 30,000 files) I transferred to Amazon EC2 the following: 1) the data from one group of 5 gamblers, 2) the play style and policy function coefficients associated with one gridpoint, 3) a generic c++ file to execute the forward simulation algorithm, and 4) a generic file that loads the starting data and coefficient estimates, runs the forward simulation algorithm, and saves the output. Within an hour the shell script revisited the Amazon EC2 instance and collected the output file. The 30,000 output files were then consolidated in preparations for the Bayesian estimation. Details of the forward simulation algorithm are available in the Appendix of the paper.

Step 3. Bayesian Estimation

After the forward simulation files were consolidated, I proceeded with the Bayesian estimation. The consolidation file is structured into an array where the rows are each row of the data, the columns are the discounted values, and the slices represent the gridpoints. The memory demands of this array are the primary reason I had to limit the estimation to 1,000 gamblers. Recall that since the parameter space is large, the Metropolis-Hastings random walk procedure is split into four partitions:

Partition	Parameters
BLOCK 1	$\{A_0, \sigma_0^2\}$
BLOCK 2	$\{\theta_0, \theta_1, \theta_2, \theta_3, r\}$
BLOCK 3	$\{\Omega\}$
BLOCK 4	$\{\Gamma\}$

Below is an outline of the homogeneous estimation procedure.

1. Initiate all parameters at "old" values.

2. For each MCMC draw:

   (a) Draw "new" BLOCK 1 learning process priors A_0 and σ_0^2 and truncate if they are outside of the grid range.

   (b) Using bi-linear interpolation, interpolate "new" discounted forward simulation values using the newly drawn learning process priors.

   (c) Draw "new" values for all other parameters: BLOCK 2, BLOCK 3, and BLOCK 4.

   (d) Process BLOCK 1 parameters:

      i. Keeping BLOCK 2, BLOCK 3, and BLOCK 4 parameters fixed at the "old" levels, compute $\alpha = \min\left[1, \frac{L^{new}}{L^{old}}\right]$ and accept the "new" BLOCK 1 parameters with probability α.

      ii. If the BLOCK 1 parameters are accepted, override the "old" values with the "new" values.

   (e) Process BLOCK 2, BLOCK 3, and BLOCK 4 parameters in a similar fashion.

3. End MCMC loop

In the homogeneous estimation, the "new" parameters are generated by a random-walk process. That is, $\theta^n = \theta^o + \varepsilon$, where $\varepsilon \sim \mathcal{N}\left(0, \sigma^2\right)$. Each parameter is drawn independently of the others. The variances of each parameter are given below:

Coefficient	Description	Random-Walk Variance
θ_0	Intercept	5e-5
A_0	Prior Mean	1e-5
σ_0^2	Prior Variance	5e-8
θ_1	Cost	1e-9
r	Risk	1e-17
θ_2	Gaming Offer	1e-7
θ_3	Room Offer	1e-7
ω_1	Weeks since last trip^1	2e-8
ω_2	Weeks since last trip^2	2e-10
ω_3	Weeks since last trip^3	4e-19
ω_4	Weeks since last trip^4	2e-28
ω_5	Weeks since last trip^5	2e-33
$\gamma_1 - \gamma_{11}$	Monthly Dummies	3e-3

For the hierarchical estimation, all variables except the $\{\Omega\}$ weeks since last trip polynomials and $\{\Gamma\}$ monthly dummy variables are estimated at the individual level. Below is an outline of the heterogeneous estimation procedure:

1. Initiate all parameters at "old" values.

2. Initiate the heterogeneity variance: Σ_θ.

3. For each MCMC draw:

   (a) Calculate the random walk step for the individual level parameters: $s^2 \times \Sigma_\theta$.

   (b) For each gambler i:

      i. Draw "new" individual level parameters $\left\{A_{0i}, \sigma_{0i}^2, \theta_{0i}, \theta_{1i}, \theta_{2i}, \theta_{3i}, r_i\right\}$ using variance Σ_θ.

      ii. Truncate A_{0i} and σ_{0i}^2 if they are outside of the grid range.

      iii. Using bi-linear interpolation, interpolate "new" discounted forward simulation values using the newly drawn learning process priors.

      iv. Process BLOCK 1 parameters for gambler i:

A. Keeping BLOCK 2, BLOCK 3, and BLOCK 4 parameters fixed at the "old" levels, compute $\alpha = \min\left[1, \frac{L_i^{new}}{L_i^{old}}\right]$ and accept the "new" BLOCK 1 parameters with probability α.

B. If the BLOCK 1 parameters are accepted, override the "old" values with the "new" values.

v. Process BLOCK 2 parameters for gambler i in a similar fashion to BLOCK 1.

(c) End gambler loop

(d) Draw "new" Δ and Σ_θ using a multivariate regression of $\Theta = Z\Delta$, where the ith row of Θ is gambler i's θ and the ith row of Z is gambler i's demographics.

(e) Draw "new" values for BLOCK 3 and BLOCK 4 parameters. As in the homogeneous estimation, each parameter is drawn independently of each other.

(f) Process BLOCK 3 parameters:

i. Keeping the BLOCK 4 parameters at the "old" levels and the BLOCK 1 and BLOCK 2 parameters at the mean level across gamblers, compute $\alpha = \min\left[1, \frac{L^{new}}{L^{old}}\right]$ and accept the "new" BLOCK 3 parameters with probability α.

ii. If the BLOCK 3 parameters are accepted, override the "old" values with the "new" values.

(g) Process BLOCK 4 parameters in a similar fashion to BLOCK 3.

4. End MCMC loop

The heterogeneous prior settings are given below:

$$\bar{\Delta} = \begin{array}{lccccccc} & A_0 & \sigma_0^2 & \theta_1 & \theta_2 & \theta_3 & \theta_4 & \theta_0 \\ \text{Intercept} & 0.397 & 0.001 & -0.002 & 0.000 & 0.005 & 0.017 & -2.860 \\ \text{Age } (/10) & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \text{Male} & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \text{Log (Distance)} & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \text{Gold LP Card} & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{array}$$

$$A = \mathrm{diag}\,(1)$$

$$\nu = 10$$

$$
\begin{aligned}
\Sigma &= \begin{bmatrix} 1e-5 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1e-8 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1e-9 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1e-17 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1e-7 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1e-7 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 5e-5 \end{bmatrix} \\
s^2 &= .2
\end{aligned}
$$

The random walk step sizes for the heterogeneous BLOCK 3 and BLOCK 4 parameters are the following:

Coefficient	Description	Random-Walk Variance
ω_1	Weeks since last trip^1	1e-9
ω_2	Weeks since last trip^2	1e-11
ω_3	Weeks since last trip^3	2e-20
ω_4	Weeks since last trip^4	1e-29
ω_5	Weeks since last trip^5	1e-34
$\gamma_1 - \gamma_{11}$	Monthly Dummies	5e-3