How Do Individuals Repay Their Debt? The Balance-Matching Heuristic

Online Appendix

John Gathergood Neale Mahoney Neil Stewart Jörg Weber

August 16, 2018

A Machine Learning Models

This section provides details of machine learning models we use to fit repayment behavior. We estimate decision tree, random forest and extreme gradient boosting. For all of these models, our target variable is the percentage of payments allocated to the high APR card in the two-card sample. We use APRs, balances, and credit limits on both cards as explanatory variables, and tune the models with cross-validation to maximize the out-of-sample power.

Decision Tree Tree-based methods partition the sample space into a series of hyper-cubes, and then fit a simple model in each partition. The decision tree is grown through iteratively partitioning nodes into two sub-nodes according to a splitting rule. In our case, the splitting criterion is to find one explanatory variable and a cut-off value that minimize the sum of squared errors in the two sub-nodes combined. In theory, the tree can have one observation in each final node, but this tree will have poor performance out-of-sample. In practice, the decision tree is grown until the reduction in squared error falls under some threshold. Then, it calculates the average percentage of payments allocated to high APR cards in each final node.

We use the r package “rpart” to fit the decision-tree model.[1] To avoid overfitting the data, we “prune” the decision tree by tuning the complexity parameter through cross-validation. The complexity parameter requires each split to achieve a gain in R-squared greater than the parameter value. We pick the complexity parameter threshold that minimizes mean square error in 5-fold cross-validation. That is, we split the sample randomly into 5 disjoint subsets. For each of these 5 subsets, we use the remaining 80 percent of the data to train the tree, and calculate the error on each 20 percent subset.[2] Appendix Figure A9 shows the estimated decision tree.

Random Forest The machine learning literature has proposed several variations on the tree model. One approach which has been found to work very well in practice is random forest (Breiman 2001). Random forest builds a large number of trees and averages their predictions. It introduces randomness into the set of explanatory variables considered when splitting each node. The algorithm first draws a number of bootstrapped samples, and grows a decision tree within each sample. At each node, it randomly selects a subset of “m” explanatory variables in the split search, and chooses the best split among those “m” variables. Lastly, it makes predictions by averaging the results from each tree.

We use the r package “randomForest” to grow a forest of 100 trees.[3] For each tree, we calculate the out-of-sample error using the rest of the data not included in the bootstrapped sample. The average prediction error over these 100 trees is minimized to fine tune “m,” the number of explanatory variables in the subset we consider in each split search. Increasing the number of trees does not significantly improve prediction accuracy.

Extreme Gradient Boosting Extreme gradient boosting and random forest are both based on a collection of tree predictors. They differ in their training algorithm. The motivation for boosting is a procedure that combines the outputs of many “weak” classifiers to produce a powerful “committee” (Friedman, Hastie and Tibshirani 2001). Instead of growing a number of

[1] See `https://cran.r-project.org/web/packages/rpart/vignettes/longintro.pdf` for a complete description of the function.

[2] See Friedman, Hastie and Tibshirani (2001) Chapter 9, for further information on tree-based methods.

[3] See `https://cran.r-project.org/web/packages/randomForest/randomForest.pdf` for a complete description of the function.

trees independently, boosting applies an additive training strategy, by adding one new tree at a time. At each step, the new decision tree puts greater weights on observations that are misclassified in the previous iteration. Finally, it averages predictions from trees at each step. This algorithm effectively gives greater influence to the more accurate tree models in the additive sequence. We use the r package "xgboost" and fine tune the number of iterations over a 5-fold cross-validation.[4] The rest of the parameters such as the learning rate are kept at their default values. Perturbation of these values does not have material impact on out-of-sample errors.[5]

B Costs of Misallocation: Extensions

In the main text, we presented the annualized interest savings from a counterfactual "steady state" where individuals optimize balances across the credit cards we observe in our data, subject to the constraint of not exceeding their credit limits. In this section, we present two extensions of these baseline calculations, focusing on the two card sample for tractability.

First, in Table A3 we present interest savings calculated using the baseline steady state approach for observations that were excluded from our baseline sample. The top row reproduces the estimates for the baseline two-card sample, also show in Panel B of Table 2. The subsequent rows show savings for observations excluded by different sample restrictions. The interest savings are roughly one-third as large for individuals who are excluded for paying the minimum on both cards and comparable for observations where individuals are revolving on one card only. The interest savings are zero for observations that are excluded because the individual pays both cards in full, has equal interest rates, or does not carry a revolving balance. These individuals are not borrowing, so there is no borrowing to optimize. Total savings, which combine positive savings for borrowers and zero savings for non-borrowers, are shown in the final row. Average interest savings for this sample are roughly one-third those in the baseline sample.

[4] See `http://cran.fhcrc.org/web/packages/xgboost/vignettes/xgboost.pdf` for a complete description of the function.

[5] For a more detailed introduction of extreme gradient boosting, see `http://xgboost.readthedocs.io/en/latest/model.html`. Friedman (2001) is the first paper that introduced the term "gradient boosting." Friedman, Hastie and Tibshirani (2001), Chapter 10 also introduces a boosting algorithm.

Second, we show interest savings from exercise where counterfactual interest costs are determined by simulating forward outcomes when individuals optimally repay their credit cards over time. This simulation is supposed to measure the gains from "learning" the optimal repayment rule, which the individual can then implement over time, and incrementally shift their balances across cards.

An important complication with applying this alternative model is that we have to take a stand on counterfactual spending behavior. Assuming their spending "stays the same" is not an option: With the updated repayments, some individuals will go over their credit limits on the lower APR card. One option is to assume that individuals allocate their spending optimally. The optimal allocation of spending is achieved by prioritizing spending to the low APR card, allocating spending only to the high APR card once the credit limit of the low APR card is reached. If you told someone how to repay optimally, a good guess is that they would also adjust their spending towards the optimal allocation (although we agree that they might not get all the way to optimal). Another option is to assume that they hold their spending fixed, unless they bump into a credit limit, in which case we can reallocate their spending to the other card. We show counterfactual interest payments under both of these assumptions for spending behavior.

We implement these calculations on our two-card sample. To capture the counterfactual where individuals learn about optimal spending and incrementally shift their balances over time, we need to observe individuals without gaps for multiple months. To create a balanced panel, we draw a sample of individuals who enter the individual $\times$ month sample restrictions in at least one month of the data and then remain in the data for at least 11 subsequent months (with those months either inside or out of the sample restrictions). This sample therefore differs from the baseline pooled sample of observations.

Table A12 shows summary statistics for interest savings. The top panel shows savings when optimizing both payments and spending, the bottom panel shows savings from optimizing payments only. Interest savings in both versions of the dynamic optimal model are weakly positive. When individuals optimize both payments and spending (Panel A), mean annualized interest savings at 12 months are close to those from the steady state calculation (£58 versus

£65) and the percentiles of the distribution are also similar. When individuals only optimize payments, balances take longer to converge to the steady state optimal allocation, and at a 12 months time horizon savings are one-third lower than savings when individuals optimize on both margins.

C Sensitivity Analysis

C.A Minimum Payment Matching

An alternative explanation for the balance-matching result is that individuals anchor their payments to minimum payment amounts. Like balances, minimum payments are prominently displayed on credit card statements (see Figure A5). If repayments are determined by a minimum-payment-matching heuristic, and minimum payments are proportional to balances, then minimum payment matching could produce the observed repayment behavior.[6]

We separately identify balance matching from minimum payment matching by "zooming in" on a subset of observations where predicted payments under balance matching and minimum payment matching are very different. This approach is better than including minimum payment matching as another heuristic in the goodness-of-fit analysis. If the balance-matching and minimum-payment-matching amounts were largely overlapping, both heuristics would have similar goodness-of-fit, even if repayments were driven by only one model of behavior.

To understand how we separately identify these two explanations, we need to provide some background on minimum payment formulas. Most minimum payment amounts are calculated as the maxim of a fixed amount and a percentage of the balance. For instance, a typical minimum payment formula might be:

$$\text{Minimum Payment} = \max\{£25, 2\% \times \text{Balance}\}.$$

Consider the following scenarios for an individual with two cards:

(i) If minimum payments are on the "fixed" part of the formula (balances greater than £1,250),

[6] Setting payments at multiples of the minimum amount (e.g., twice the minimum on each card) would also produce the observed repayment behavior.

and the percentages are identical (2 percent for both cards), then the balance-matching and minimum-payment-matching payments will be almost perfectly correlated.[7]

(ii) If the percentages differ, then balance-matching and minimum-payment-matching payments will be correlated, but to a lesser extent.

(iii) If minimum payments are on the "percentage" part of the formula (balances less than £1,250), then the balance-matching allocation will not be correlated with the minimum-payment-matching allocation.

Hence, focusing on observations that have different percentages in the minimum payment rule (scenario ii) and where the fixed payment binds (scenario iii) allows us to separately identify these mechanisms.

Figure A10 shows binned-scatter plots of actual and predicted payments on the high interest rate card under the balance-matching heuristic (left column) and minimum-payment-matching heuristic (right column). The top row shows this relationship where both cards have the same percentage (scenario i), the middle row shows this relationship when the percentages are different (scenario ii), and the bottom row shows this relationship when both cards are on the fixed part of the formula (scenario iii). The correlations between these different measures are shown in Table A13.

In the same percentage sample, the balance-matching and the minimum-payment-matching payments are near-perfectly correlated ($\rho = 0.96$). As a result, the correlation between actual and balance-matching payments ($\rho = 0.63$) is nearly identical to the correlation between actual and minimum-payment-matching payments ($\rho = 0.61$). In the different percentages sample, the balance-matching and the minimum-payment-matching payments are more weakly correlated ($\rho = 0.86$), and the correlation between actual and balance-matching payments ($\rho = 0.41$) is stronger than the correlation between actual and minimum-payment-matching payments ($\rho = 0.28$). In the fixed sample, there is a much weaker correlation between the balance-matching payments and the minimum-payment-matching payments ($\rho = 0.56$), and the correlation between actual and balance-matching payments ($\rho = 0.50$) is substantially stronger than the correlation between actual and minimum-payment-matching payments

[7] The correlation is not perfect because minimum payment amounts may include fees incurred during the cycle, such as cash advance fees or foreign currency exchange fees.

($\rho = 0.23$).

It follows that observed repayment behavior is driven by balance matching and not by individuals setting payments in relationship to minimum payments. The correlation between actual and balance-matching payments is not affected by whether minimum-payment-matching payments are correlated with the balance-matching payment amount. On the other hand, the correlation between actual and minimum-payment-matching payment seems highly sensitive to whether the balance-matching payments are correlated with the minimum-payment-matching amount. We note that while minimum payments do not seem to be driving our findings, our analysis does *not* imply that minimum payments are irrelevant for repayment behavior. Indeed, while not directly comparable, our finding of a modest correlation between actual and minimum payments matching repayments is consistent with Keys and Wang (2018), who estimate that 9 percent to 20 percent of account-holders anchor their repayments to minimum payment amounts.

C.B Autopay

Another factor that might affect repayment behavior is whether the individual uses automatic payment ("autopay"). In the completely unrestricted two-card sample (including individuals with no revolving debt on either card), autopay is used on 23.9 percent of account $\times$ months. Although individuals are allowed to set automatic payments at a fixed amount or a fixed percentage of the balance, individuals typically set automatic payments at either the minimum due or the full balance. Conditional on using autopay, 30.3 percent pay the minimum and 42.2 percent pay the full amount. Since we drop individuals who make the minimum or full payment on both their cards (see Section 2), autopay is used on only 17.4 percent of account $\times$ months in the baseline sample. Thus, the main results predominately reflect behavior when individuals do *not* use autopay and make active repayment decisions.

Appendix Figure A11 plots repayment behavior for observations where individuals use autopay on both cards (left column, 11 percent of observations) and do not use autopay on either card (right column, 77 percent of observations).[8] The top row shows the distributions

[8] The propensity to use autopay is highly correlated within individuals across cards. In the two-card sample, 68.2 percent of individuals use autopay on the high APR card also use it on the low APR card, and 74.9 percent of

of actual and optimal payments, the middle row shows the distribution of actual and optimal payments in excess of the minimum, and the bottom row shows the joint distribution of actual and balance-matching payments. While average misallocated repayments are lower in the autopay sample than the non-autopay sample (7.3 percent versus 23.2 percent), misallocated repayments in excess of the minimum are similar in both samples (45.5 percent versus 45.7 percent). The reason that misallocated payments are smaller (and misallocated excess payments are the same) is that the autopay sample has lower monthly repayments and, therefore, the scope for misallocating payments is lower.[9] Summary statistics for actual and excess payments by autopay status are shown in Appendix Table A14.

Appendix Table A15 and Table A16 show our standard measures of model performance by whether individuals use autopay on both cards and do not use autopay on either card.[10] In particular, Appendix Table A15 shows our measures of goodness-of-fit (root mean square error, mean absolute error, Pearson's correlation) for uniformly distributed repayments, optimal repayments, and balance-matching repayments separately for the autopay and non-autopay samples. Appendix Table A16 shows the results of horse-race analysis that compares uniformly distributed versus balance-matching payments, and balance-matching versus optimal repayments, separately for the autopay and non-autopay samples. While the exact results vary, the optimal model performs poorly and the balance-matching model performs well across all of these different measures of model performance in both the autopay and non-autopay sample. Thus, we conclude that our results are not particularly sensitive to whether individuals use autopay.

In summary, autopay is rare in our baseline sample, and our main results predominately reflect repayments by individuals who do not use autopay and necessarily make active repayment decisions each month. However, when individuals use autopay, their propensity to misallocate and to follow a balance-matching rule is similar to that in the non-autopay sample, suggesting that our results are robust across these somewhat different choice environments.

individuals who do not use autopay on the high APR card do not use it on the low APR card.

[9] Specifically, while balances are slightly higher in the autopay sample (£6,900 versus £5,800), repayments are substantially lower (£200 versus £510).

[10] Results are shown using the 20 percent hold-out sample.

Figure A1: Actual and Optimal *Excess* Payments

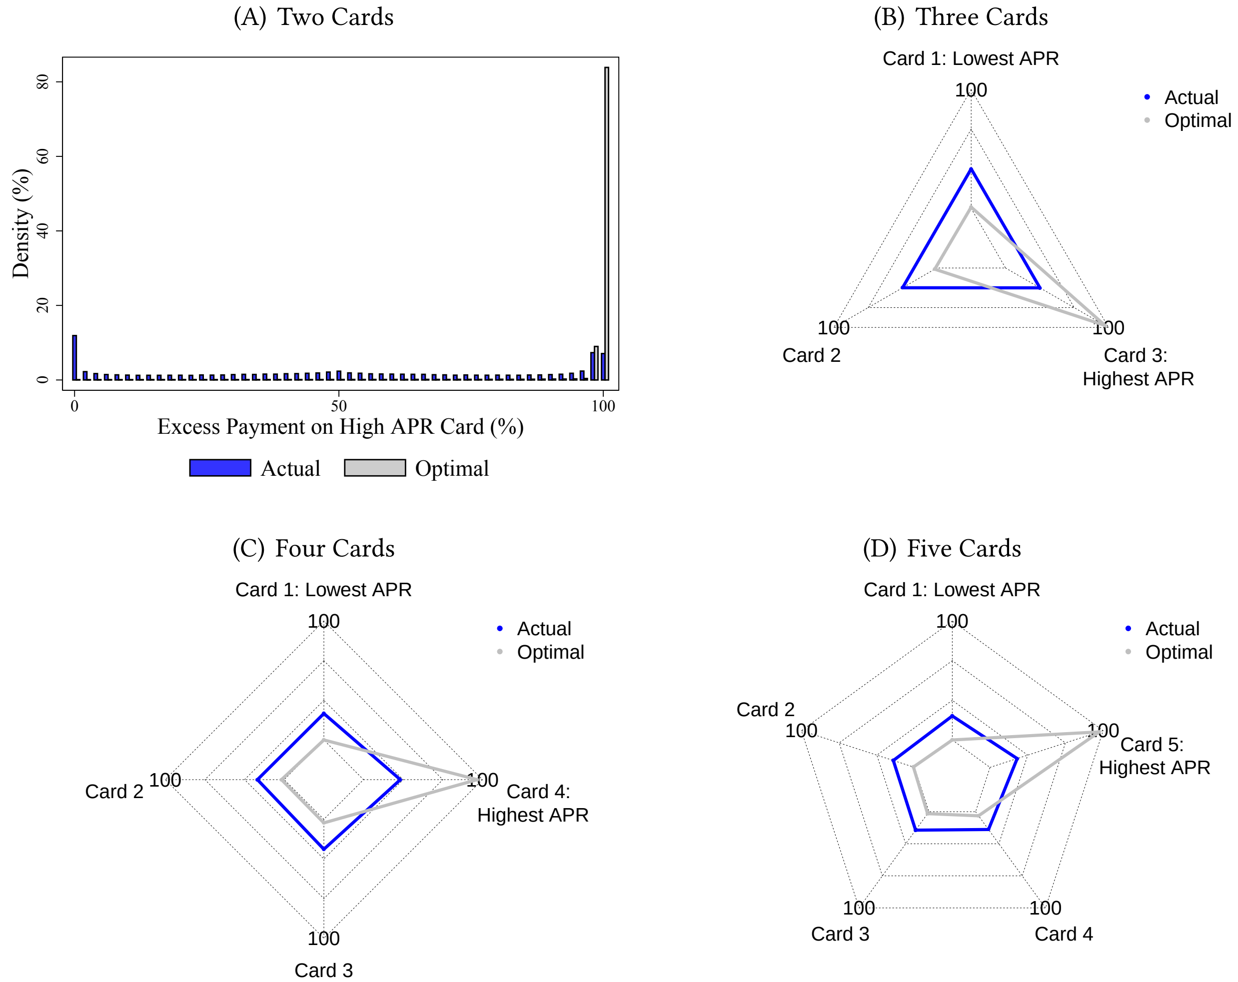

Note: Panel A shows the distribution of actual and optimal excess payments on the high interest rate card in the two-card sample. Panels B to D show radar plots of mean actual and optimal excess payments in the samples of individuals with 3 to 5 cards. Excess payments are calculated as the percentage of payments on a given card after subtracting out repayments needed to pay the minimum amounts due. In the radar plots, cards are ordered clockwise from the highest to the lowest APR (starting at the first node clockwise from noon). All samples are restricted to individual × months in which individuals face an economically meaningful allocative decision. See Section 2.2 for more details on the sample construction.

Figure A2: Misallocated *Excess* Payments by Economics Stakes

(A) Misallocated vs. Difference in APR

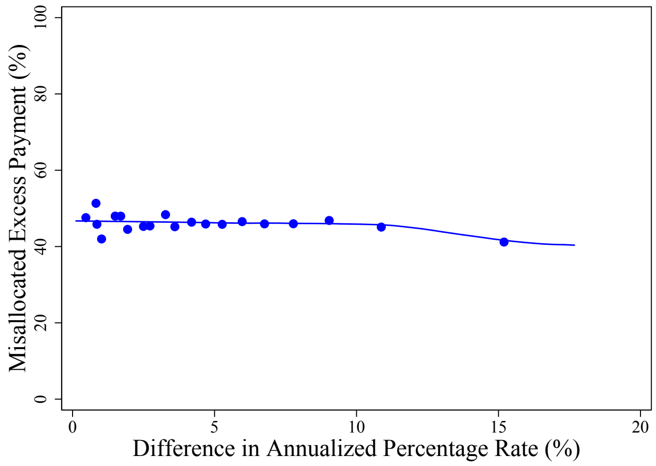

(B) Misallocated vs. Total Payments

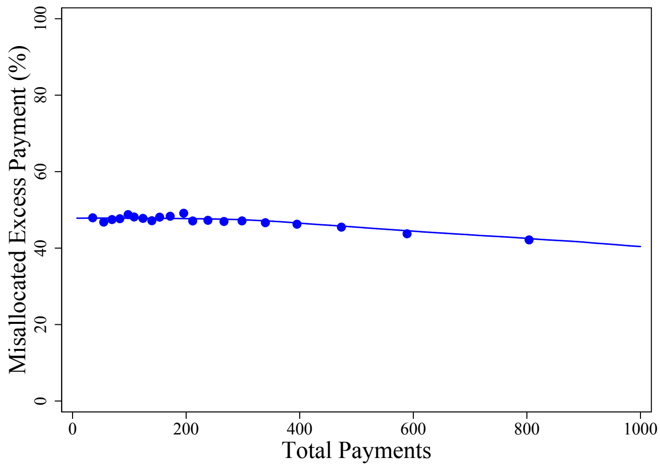

(C) Misallocated vs. Δ APR × Total Payment

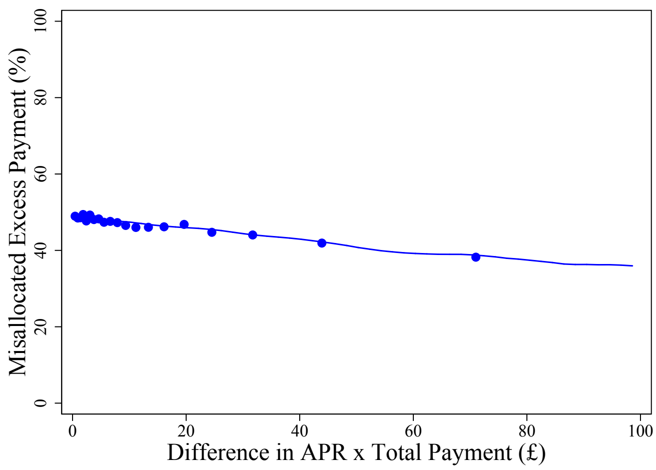

Note: Figure shows binned-scatter plots (with 20 equally sized bins) of misallocated payments in excess of the minimum payment against the difference in APR across cards (Plot A), the total value of payments within the month in pounds (Plot B) and the difference in APR multiplied by the total value of payments within the month (Plot C). Local polynomial lines of best fit, based on the non-binned data, are also shown. The two-card sample is restricted to individual × months in which individuals face an economically meaningful allocative decision. See Section 2.2 for details on the sample construction.

Figure A3: Misallocated *Excess* Payments by Card Age and Difference in Due Dates

(A) Misallocated vs. Age of High-APR Card

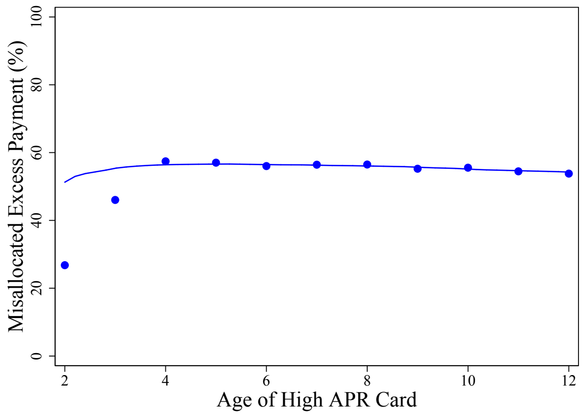

(B) Excess Misallocated Payments vs. Diff. Due Dates

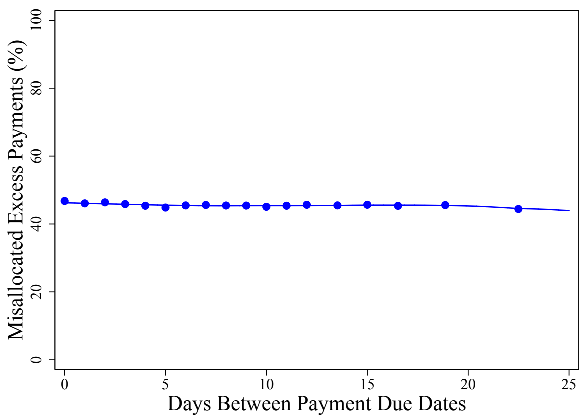

Note: Figure shows binned-scatter plots (with 20 equally sized bins) of misallocated payments in excess of the minimum payment against the difference in payment due dates (Plot A) and age of the high APR card (Plot B). Local polynomial lines of best fit, based on the non-binned data, are also shown. The two-card sample is restricted to individual × months in which individuals face an economically meaningful allocative decision. See Section 2.2 for details on the sample construction.

Figure A4: Histogram of Difference in Due Dates

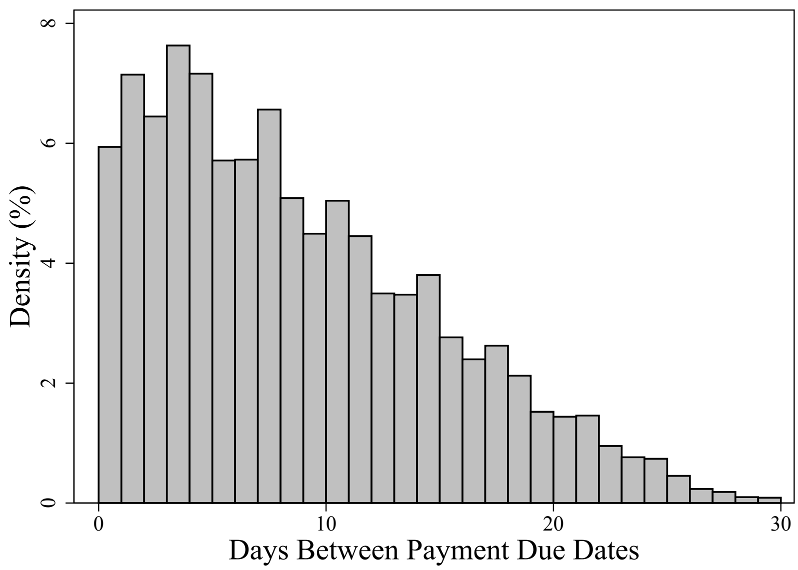

Note: Figure shows the distribution of the absolute difference in due dates. The two-card sample is restricted to individual × months in which individuals face an economically meaningful allocative decision. See Section 2.2 for more details on the sample construction.

Figure A5: Example Credit Card Statement

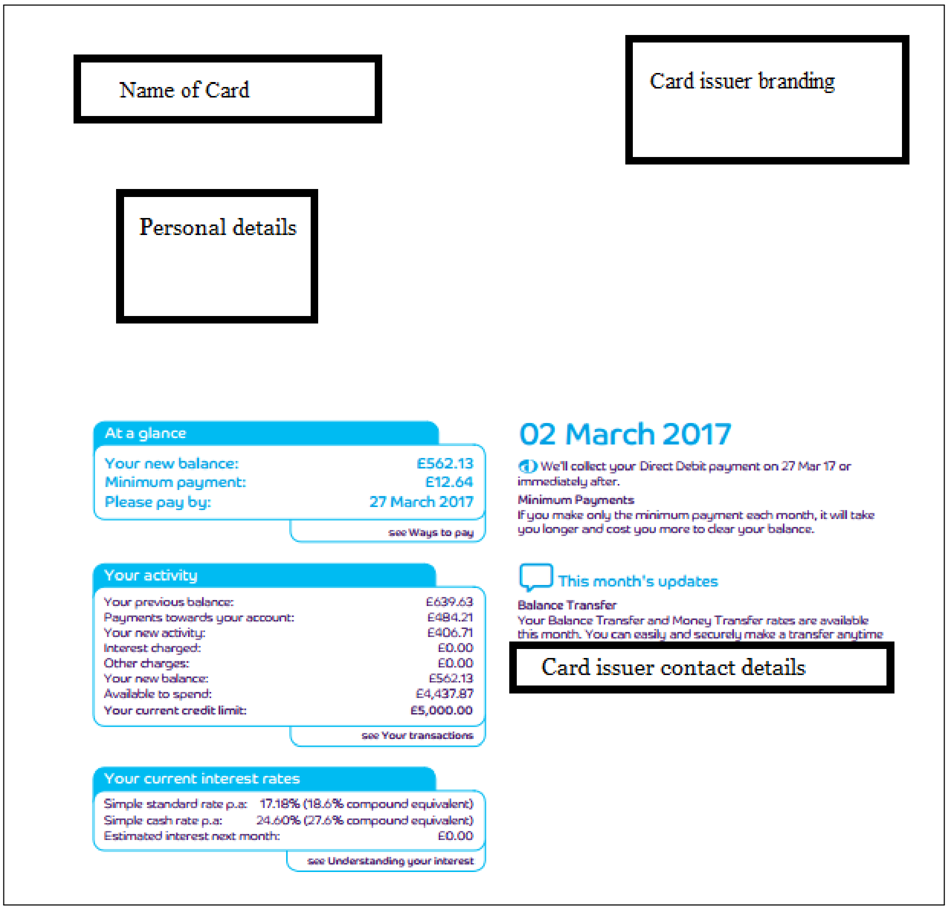

Note: The figure shows an extract of one of the authors' credit card statements, with card issuer branding, contact details and card holder personal identifying information obscured.

Figure A6: Balance Matching

(A) Round Number Payment Sample

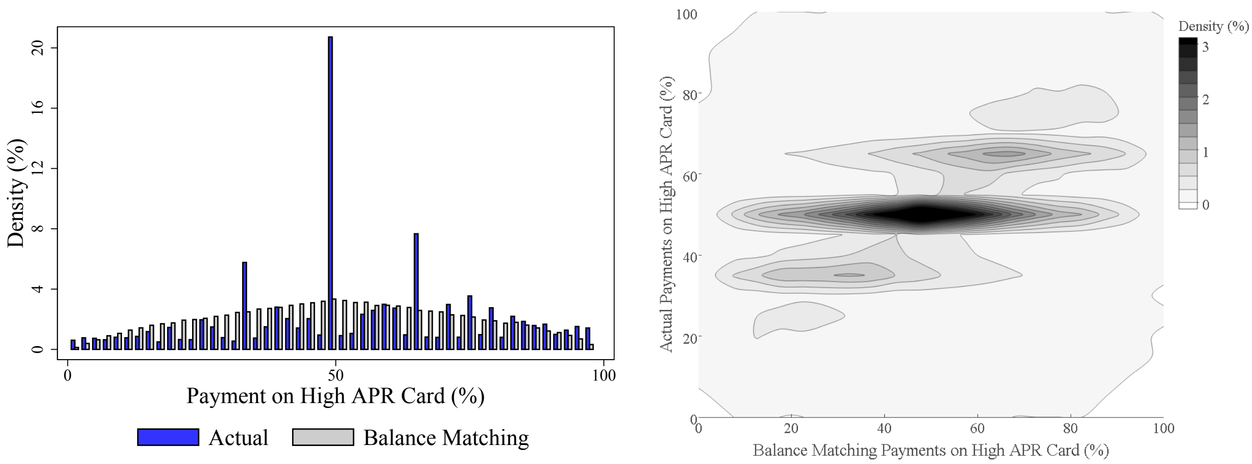

Note: Left panels shows the distribution of actual and balance-matching payments on the high APR card. Sample restricted to round number payments (multiples of £50) Round and non-round samples are defined by repayments on the high APR card. See Footnote 35 for details. The two-card sample is restricted to individual × months in which individuals face an economically meaningful allocative decision. See Section 2.2 for details on the sample construction.

Figure A7: Goodness-of-Fit for Different Models, Round and Non-Round Number Samples

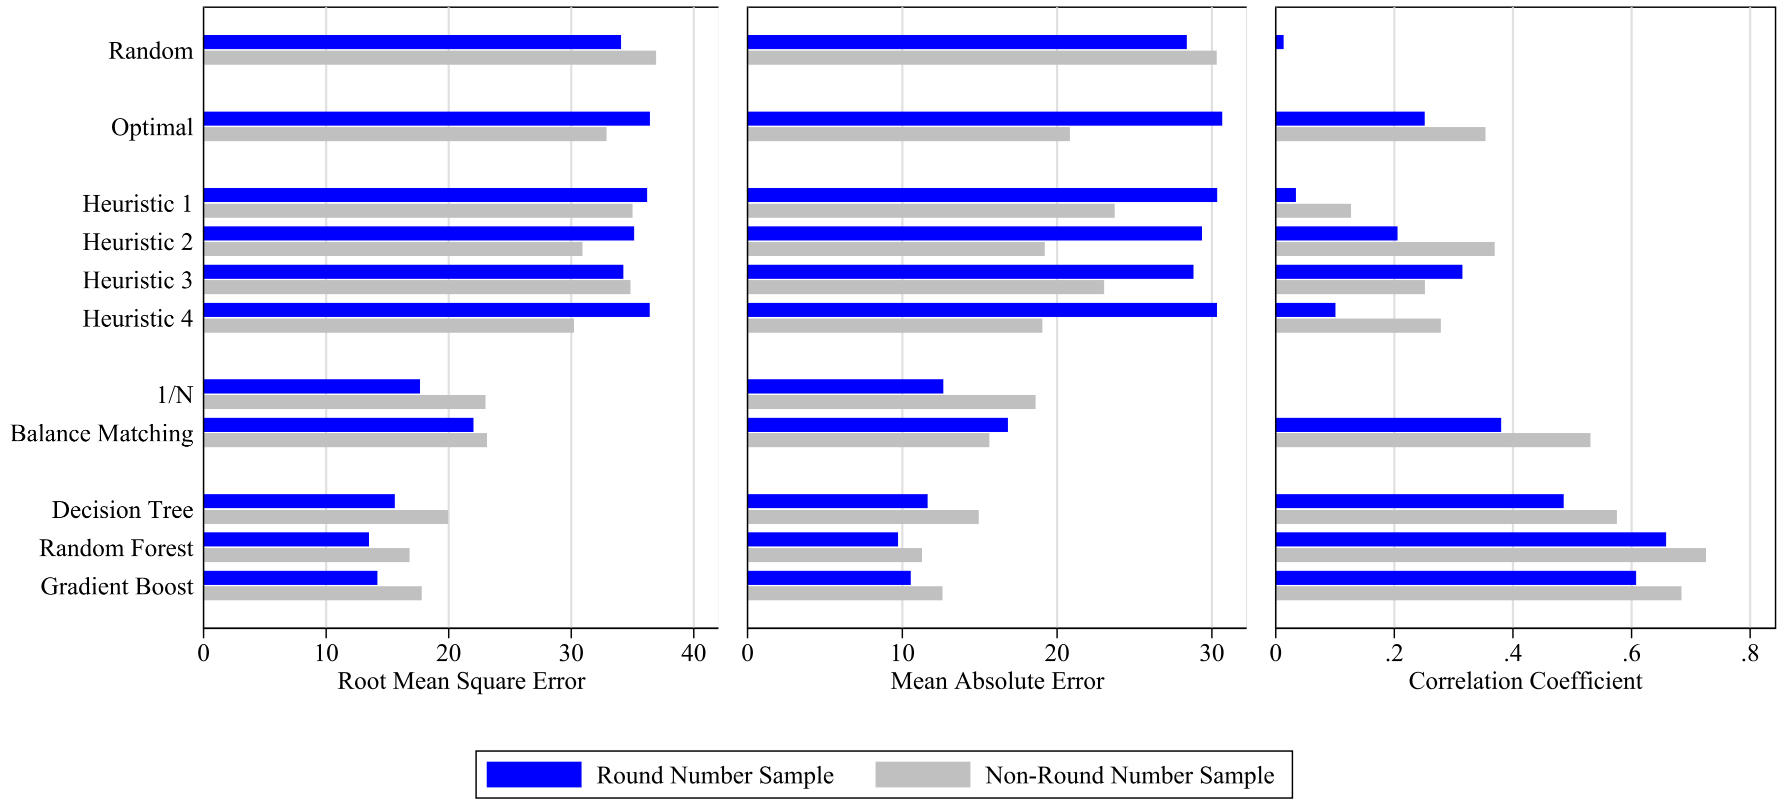

Note: Goodness-of-fit for different models of the percentage of payments on the high APR card. The left panel shows the Root Mean Square Error (RMSE), the middle panel shows the Mean Absolute Error (MAE), and right panel shows the Pearson Correlation Coefficient, which can also be interpreted as the square root of the R-squared. The round number sample restricts to observations where individuals make round number payments (multiples of £50), and the non-round number sample restricts to observations where individuals make non-round payment amounts (not multiples of £50). Random has repayments on the high APR card randomly drawn from a uniform distribution with support on the 0 percent to 100 percent interval. Optimal is pay minimum required payment on all of their cards, repay as much as possible on the card with the highest interest rate, and only allocate further payments to the lower interest rate cards if they are able to pay off the highest interest rate card in full. Heuristic 1 is repay the card with the lowest capacity. Heuristic 2 is repay the card with highest capacity. Heuristic 3 is repay the card with the highest balance. Heuristic 4 is repay the card with the lowest balance ("debt snowball method"). Balance matching is match the share of repayments on each card to the share of balances on each card. Decision Tree, Random Forest, and Gradient Boost are machine learning models that predict the share of repayments on the high APR card using these methods. Round and non-round samples are defined by repayments on the high APR card. See Footnote 35 for details. The two-card sample is restricted to individual × months in which individuals face an economically meaningful allocative decision. Goodness of fit is calculated using the 20 percent hold-out sample. See Section 2.2 for more details on the sample construction.

Figure A8: Actual and Predicted Payments Under Alternative Repayment Heuristics

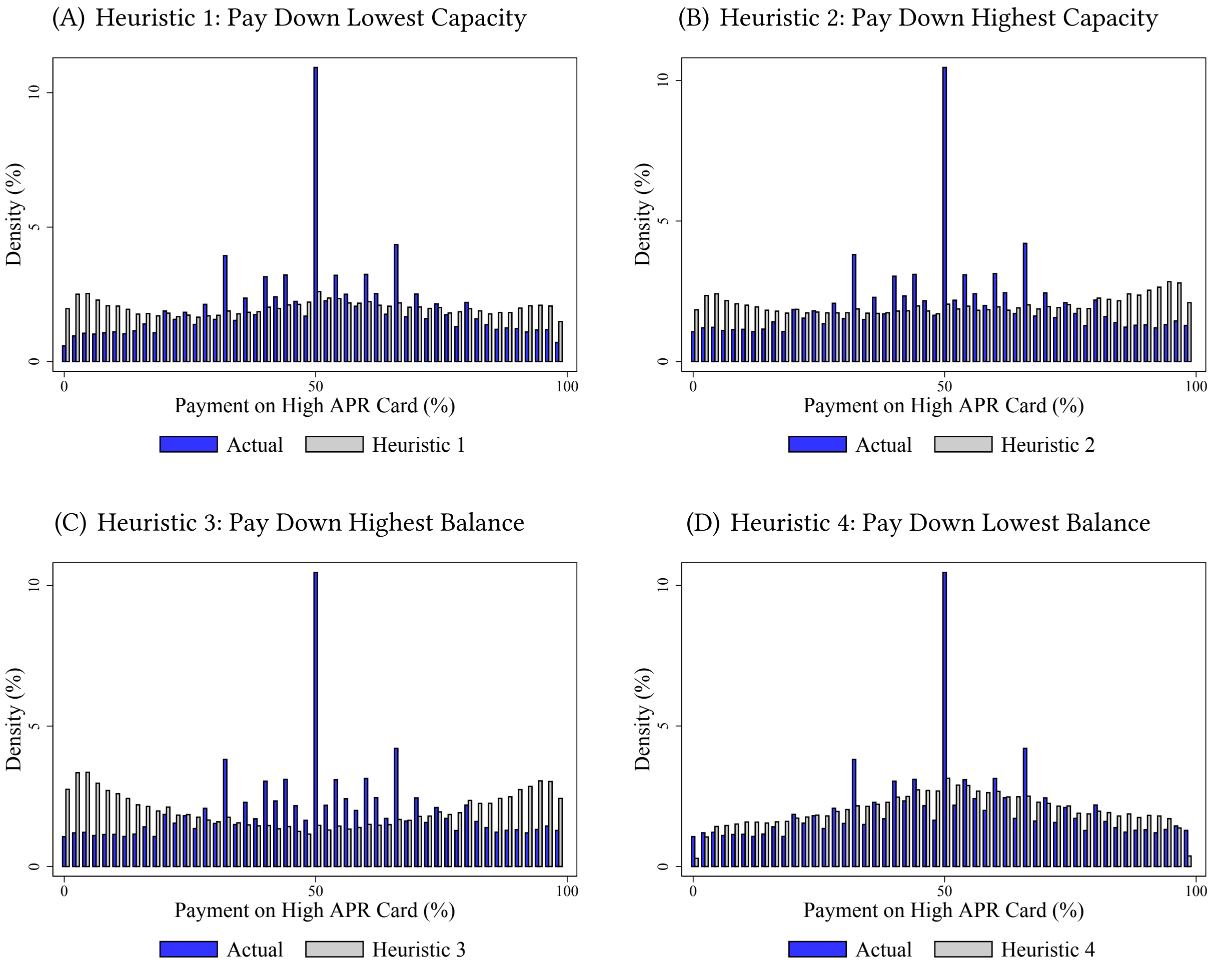

Note: Figures show the distributions of actual payments and predict payments under the alternative repayment heuristics. Heuristic 1 is repay the card with the lowest capacity. Heuristic 2 is repay the card with highest capacity. Heuristic 3 is repay the card with the highest balance. Heuristic 4 is repay the card with the lowest balance ("debt snowball method"). The two-card sample is restricted to individual × months in which individuals face an economically meaningful allocative decision. See Section 2.2 for more details on the sample construction.

Figure A9: High APR Card Payment Decision Tree

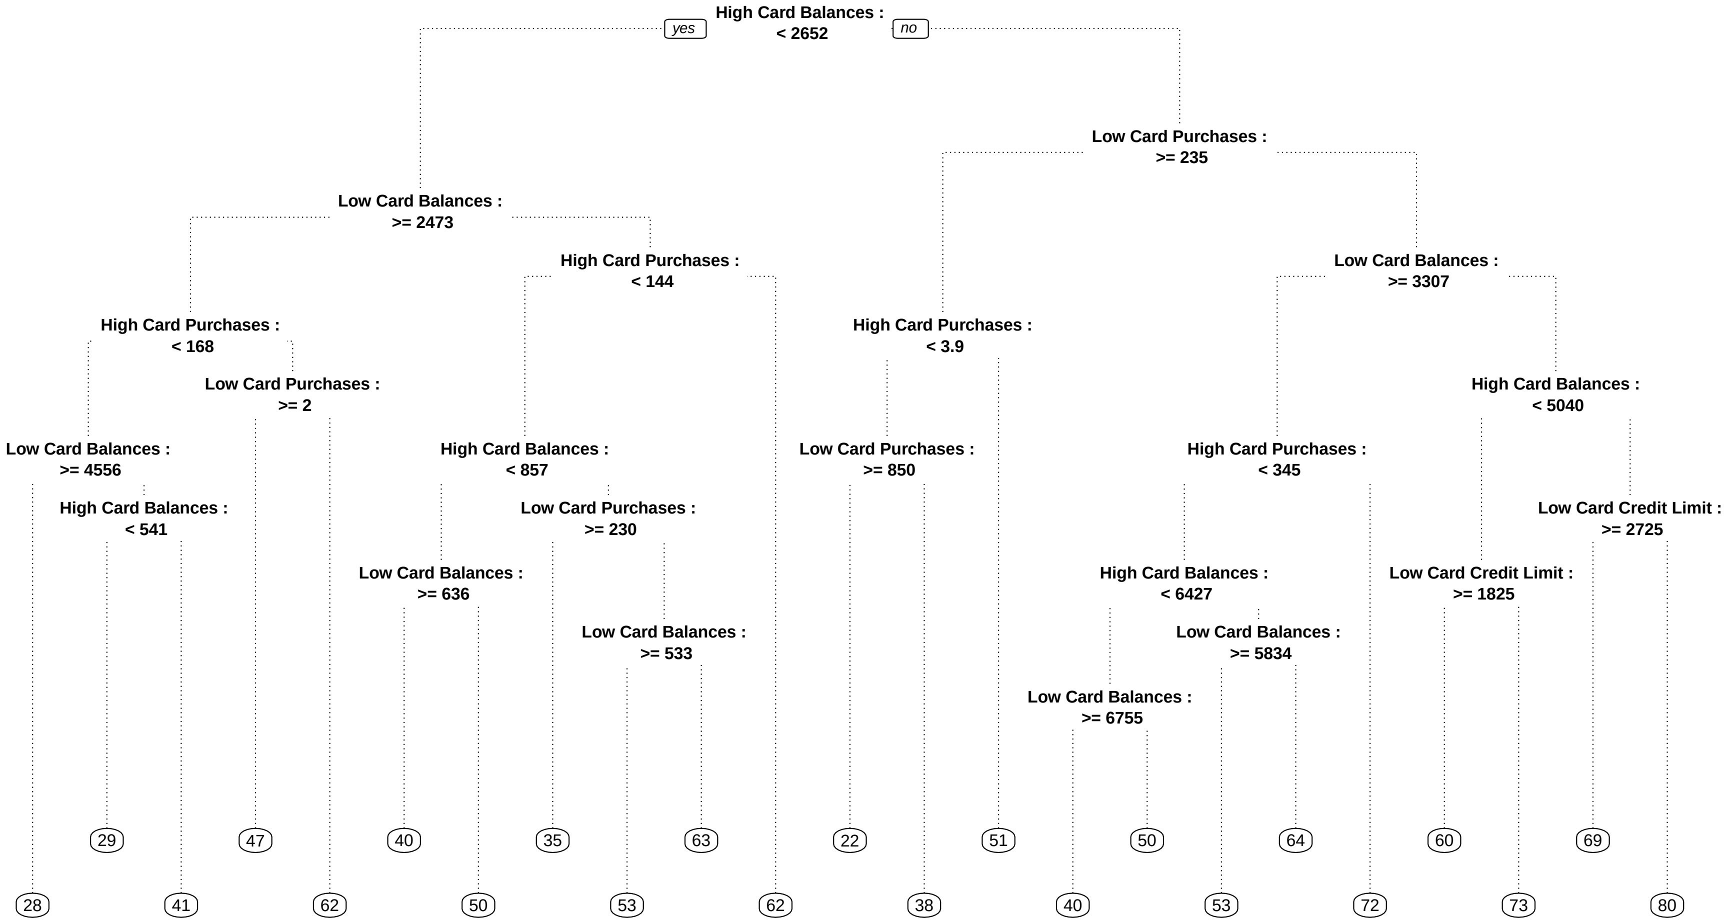

Note: Figure shows the decision (regression) tree for high APR card repayment. Top row is tree root. Nodes show the variable and split value at each branch. Bottom rows show predicted values at the end of each branch.

Figure A10: Balance Matching and Minimum Payment Matching in the Percentage and Fixed Payment Samples

(A) Same Percentage Sample

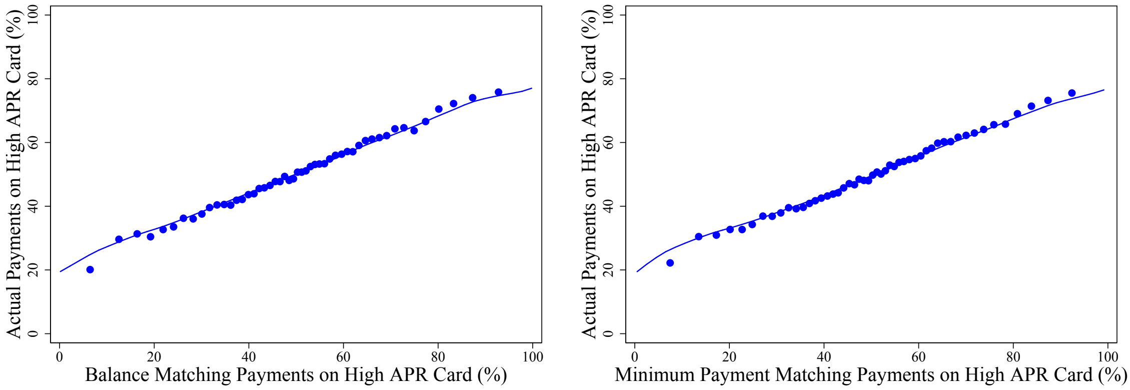

(B) Different Percentage Sample

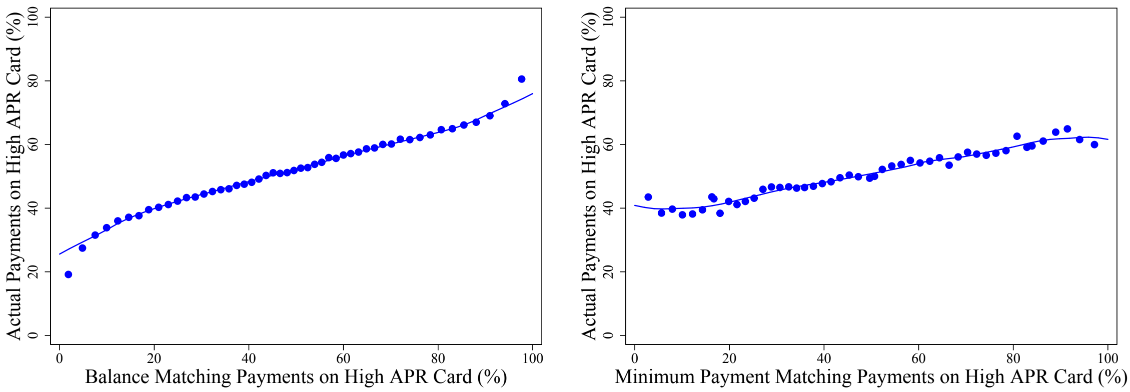

(C) Fixed Payment Sample

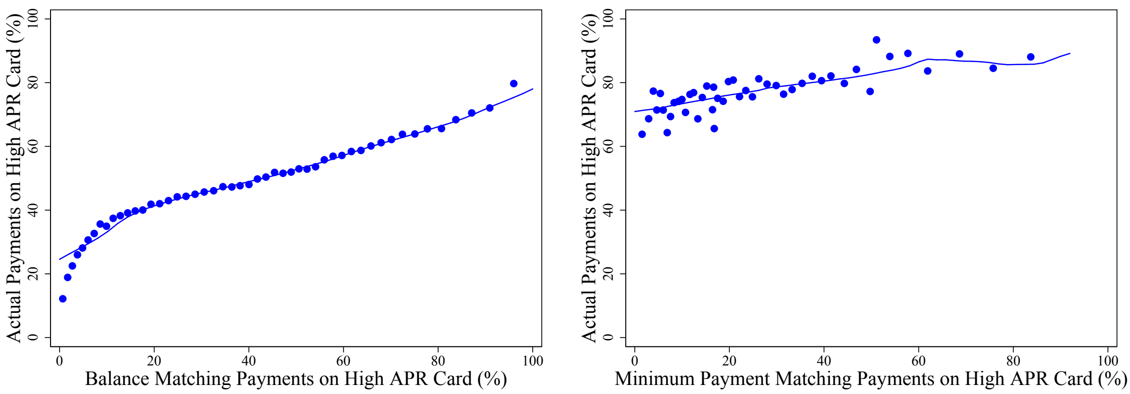

Note: Panels show binned-scatter plots of the actual percentage of monthly payment allocated to the high APR card (y-axis) and the percentage of total monthly payment allocated to the high APR card under the balance-matching heuristic (x-axis, left column) and minimum-payment-matching heuristics (x-axis, right column). "Same Percentage Sample" focuses on account × months where the balance-matching and minimum-payment-matching payments are near-perfectly correlated ($\rho = 0.96$). "Different Percentage Sample" focuses on account × months where the balance-matching and minimum-payment-matching payments are less strongly correlated ($\rho = 0.86$). "Fixed Payment Sample" focuses of account × months where the balance-matching and minimum-payment-matching payments have the weakest correlation ($\rho = 0.56$).

Figure A11: Actual, Optimal and Balance Matching Payments for Autopay (Left Column, 11 Percent of Observations) and Non-Autopay (Right Column, 77 Percent of Observations) Samples

(A) Actual vs. Optimal Payments

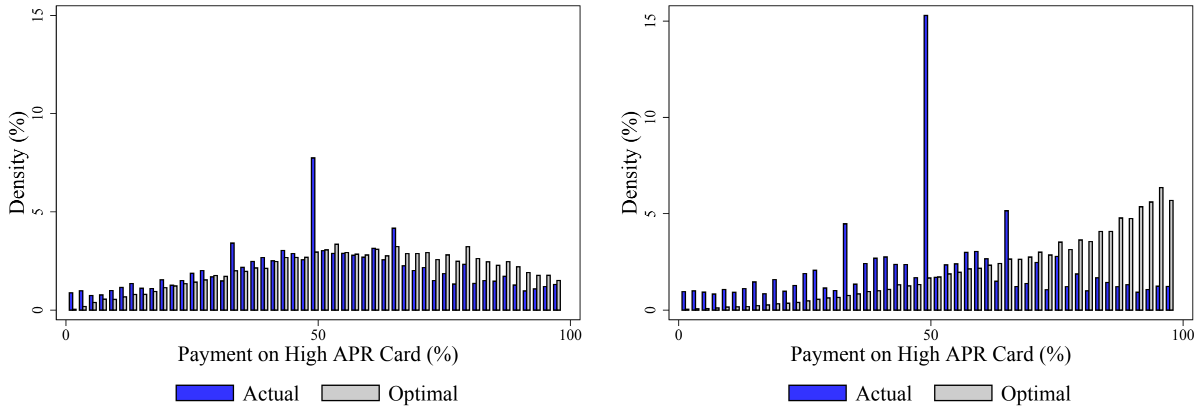

(B) Actual vs. Optimal Excess Payments

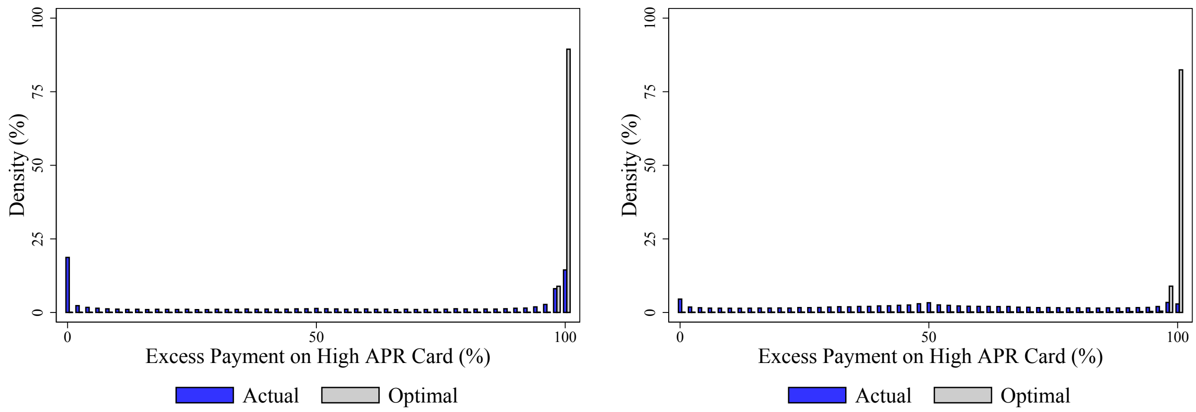

(C) Actual vs. Balance Matching Payments

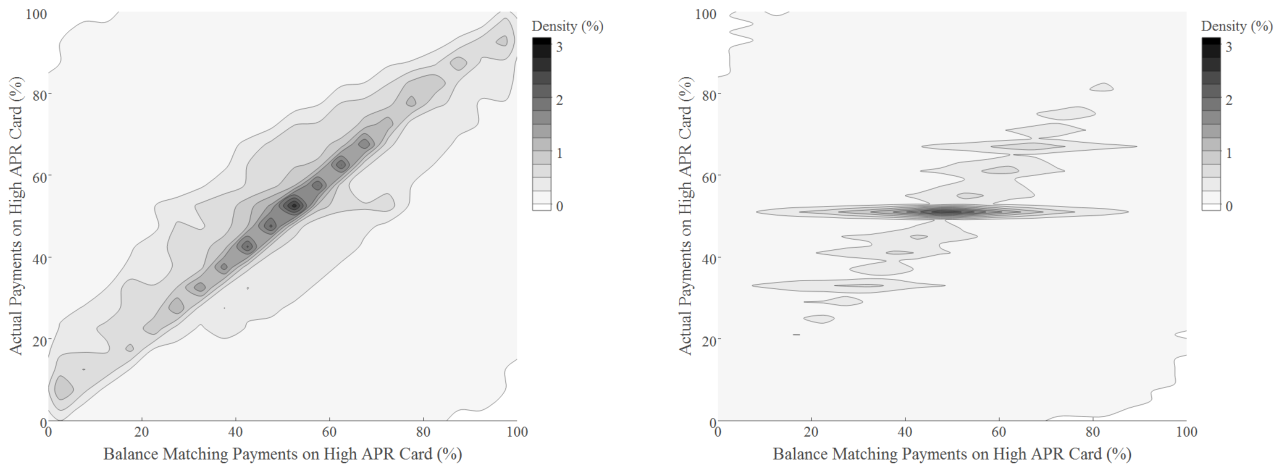

Note: Panel A shows the distribution of actual and optimal excess payments on the high APR card in the two-card sample. Panel B shows the distribution of actual and optimal excess payments on the high APR card in the two-card sample. Excess payments are calculated as the percentage of payments on a given card after subtracting out repayments needed to pay the minimum amounts due. Panel C shows the joint distribution of actual and balance matching payments on the high APR card. The autopay sample is defined as observations where individuals make automatic payments on both cards. The non-autopay sample is defined as observations where individuals do not make automatic payments on either card. All samples are restricted to individual × months in which individuals face an economically meaningful allocative decision. See Section 2.2 for more details on the sample construction.

Table A1: Sample Restrictions

	(1) Unique Individuals		(2) Aggregate Revolving Debt	
	Count	Percent	£s	Percent
Unrestricted Sample	174,686	100.00%	301,182,890	100.00%
Drop if Equal Interest Rates	2,845	1.63%	6,293,817	2.09%
Drop if Pays Full on Both	10,782	6.17%	18,239,430	6.06%
Drop if Pays Min on Both	48,263	27.63%	50,590,569	16.80%
Baseline Sample	112,796	64.57%	226,059,074	75.06%

Note: Table shows the effect of the sample restrictions on the number and percentage of unique individuals and aggregate debt in the two-card sample. Since observations may be excluded by multiple criteria, the order in which the restrictions are applied matters, and the values in the table should be thought about as the incremental effect of the different restrictions.

Table A2: Actual and Optimal *Excess* Payments on the High APR Card

	Mean	Std. Dev.	*Percentiles* 10th	25th	50th	75th	90th
i) As a Percent Total Monthly Payment							
Actual Excess Payment (Percent)	51.51	34.75	0.89	19.92	51.31	84.91	99.82
Optimal Excess Payment (Percent)	97.08	12.93	100.00	100.00	100.00	100.00	100.00
Difference (Percent)	45.56	35.05	0.00	11.40	45.34	75.70	98.39
ii) Payments in £							
Actual Excess Payment (£)	196.52	729.43	0.23	2.32	22.70	88.79	350.19
Optimal Excess Payment (£)	314.06	843.53	1.91	14.40	66.51	223.00	737.54
Difference (£)	117.54	422.14	0.00	1.00	17.80	75.00	237.47

Note: Summary statistics for actual and optimal excess payments on the high APR card. Excess payments are calculated as the percentage of payments on a given card after subtracting out repayments needed to pay the minimum amounts due. The top panel shows values as a percentage of total excess payments on both cards in that month. The bottom panel shows values in £s. The two-card sample is restricted to individual × months in which individuals face an economically meaningful allocative decision. See Section 2.2 for details.

Table A3: Annualized Interest Savings in the Unrestricted Sample

	Individuals	Revolving Debt	Interest Savings in £			
	(N)	£	Mean	Std. Dev.	75th Pctile	90th Pctile
Baseline sample	112,796	226,059,074	64.19	111.01	68.71	166.97
Pays min on both	48,263	50,590,569	21.93	49.16	17.51	49.28
Revolving on one only	12,046	29,312,832	66.33	119.92	68.63	170.16
Pays full on both	10,782	18,239,430	0	0	0	0
Equal interest rates	2,845	6,293,817	0	0	0	0
No balance on either card	47,655	0	0	0	0	0
Total	234,387	330,495,722	24.65	48.63	23.26	60.15

Note: Table shows summary statistics for annualized interest savings from a counterfactual "steady state" where individuals optimize balances across the credit cards we observe in our data, subject to the constraint of not exceeding their credit limits. See Section 2.2 for more details on the sample construction. The exchange rate was £1 = $1.32 at the midpoint of our sample period.

Table A4: Actual and Optimal *Excess* Payments on High APR Card by Quintiles of Economic Stakes and Card Age

	Quintiles				
	1	2	3	4	5
i) Difference in APR					
Difference in APR (Percent)	0.81	2.03	3.78	6.44	14.98
Actual Excess Payment (Percent)	50.09	50.45	51.20	51.45	54.38
Optimal Excess Payment (Percent)	97.00	96.96	97.36	97.66	96.40
Difference (Percent)	46.91	46.51	46.16	46.21	42.01
ii) Total Payment					
Total Payment (£)	63.82	125.98	204.23	350.22	1658.91
Actual Excess Payment (Percent)	51.98	51.14	50.78	51.37	52.28
Optimal Excess Payment (Percent)	99.51	99.08	98.87	97.86	90.32
Difference (Percent)	47.54	47.94	48.09	46.48	38.04
iii) Financial Stakes					
Financial Stakes (£)	1.20	3.70	8.08	17.50	101.91
Actual Excess Payment (Percent)	50.41	50.49	51.35	51.51	53.80
Optimal Excess Payment (Percent)	99.18	98.85	98.30	97.55	91.52
Difference (Percent)	48.77	48.36	46.94	46.04	37.71
iv) Difference in Due Dates					
Difference in Due Dates (Days)	1.03	3.48	6.48	10.83	17.93
Actual Excess Payment (Percent)	50.80	51.50	51.69	51.68	51.81
Optimal Excess Payment (Percent)	97.19	97.14	97.04	97.01	97.04
Difference (Percent)	46.39	45.63	45.35	45.33	45.23
v) Age of High APR Card					
Age of High APR Card (Months)	3.82	5.57	7.50	9.50	11.48
Actual Excess Payment (Percent)	42.24	42.10	42.20	42.73	44.04
Optimal Excess Payment (Percent)	97.50	98.56	98.66	98.12	98.19
Difference (Percent)	55.26	56.46	56.46	55.40	54.15

Note: Summary statistics for actual and optimal excess payments on the high APR card by quintiles of economic stakes and card age. Excess payments are calculated as the percentage of payments on a given card after subtracting out repayments needed to pay the minimum amounts due. Cells report mean values within the quintile. The two-card sample is restricted to individual × months in which individuals face an economically meaningful allocative decision. See Section 2.2 for more details on the sample construction. The exchange rate was £1 = $1.32 at the midpoint of our sample period.

Table A5: Annualized Interest Savings Under Different Repayment Rules

			Percentiles				
	Mean	Std. Dev.	10th	25th	50th	75th	90th
Annualized Savings in £							
Optimal Payment	64.82	115.33	2.46	7.80	24.78	70.39	167.41
Balance Matching	−4.59	89.29	−61.82	−14.91	−0.32	12.43	49.06
Heuristic 1	15.06	112.66	−50.98	−12.86	0.42	21.32	90.99
Heuristic 2	6.06	94.77	−51.09	−12.61	0.68	19.93	72.03
Heuristic 3	−4.59	89.29	−61.82	−14.91	−0.32	12.43	49.06
Heuristic 4	−1.68	103.18	−70.42	−16.85	0.60	20.47	69.62

Note: The two-card sample is restricted to individual × months in which individuals face an economically meaningful allocative decision. See Section 2.2 for more details on the sample construction.

Table A6: Goodness-of-Fit for Different Models

	(1) RMSE	(2) MAE	(3) Corr
i) Main Models			
Uniform Draw (0,100)	36.59	30.05	-0.00
	(0.08)	(0.07)	(0.00)
Optimal	35.09	25.38	0.31
	(0.12)	(0.11)	(0.00)
1/N	23.00	18.19	
	(0.06)	(0.06)	
Balance Matching	23.89	17.07	0.47
	(0.08)	(0.06)	(0.00)
ii) Alternative Heuristics			
Heuristic 1 (Pay Down Lowest Capacity)	36.46	27.28	0.08
	(0.12)	(0.11)	(0.01)
Heuristic 2 (Pay Down Highest Capacity)	33.52	23.88	0.29
	(0.13)	(0.12)	(0.01)
Heuristic 3 (Pay Down Highest Balance)	35.29	25.94	0.27
	(0.12)	(0.10)	(0.01)
Heuristic 4 (Pay Down Lowest Balance)	34.20	24.68	0.10
	(0.13)	(0.12)	(0.01)
iii) Machine Learning Models			
Decision Tree	19.42	15.03	0.53
	(0.07)	(0.05)	(0.00)
Random Forest	16.24	11.63	0.71
	(0.07)	(0.05)	(0.00)
XGBoost	17.51	13.17	0.65
	(0.07)	(0.05)	(0.00)

Note: Goodness-of-fit for different models of the percentage of payments on the high APR card. The first column shows the Root Mean Square Error (RMSE), the second column shows the Mean Absolute Error (MAE), and third column shows the Pearson Correlation Coefficient, which can also be interpreted as the square root of the R-squared. The two-card sample is restricted to individual $\times$ months in which individuals face an economically meaningful allocative decision. Goodness of fit is calculated using the 20 percent hold-out sample and standard errors are constructed by the bootstrap method. See Section 2.2 for details.

Table A7: Goodness-of-Fit for Different Models, Round Number and Non-Round Number Payment Samples

	Round Number Sample			Non-Round Number Sample		
	(1) RMSE	(2) MAE	(3) Corr	(4) RMSE	(5) MAE	(6) Corr
i) Main Models						
Uniform Draw (0,100)	34.04	28.36	-0.01	36.90	30.30	-0.00
	(0.19)	(0.19)	(0.01)	(0.11)	(0.10)	(0.00)
Optimal	36.40	30.65	0.25	32.86	20.81	0.35
	(0.18)	(0.17)	(0.01)	(0.24)	(0.20)	(0.01)
1/N	17.64	12.63	-0.01	22.99	18.60	-0.00
	(0.10)	(0.10)	(0.01)	(0.06)	(0.07)	(0.00)
Balance Matching	22.00	16.81	0.38	23.11	15.61	0.53
	(0.14)	(0.11)	(0.01)	(0.13)	(0.09)	(0.00)
ii) Alternative Heuristics						
Heuristic 1 (Pay Down Lowest Capacity)	36.17	30.33	0.03	34.98	23.71	0.13
	(0.19)	(0.17)	(0.01)	(0.21)	(0.18)	(0.01)
Heuristic 2 (Pay Down Highest Capacity)	35.11	29.34	0.21	30.90	19.19	0.37
	(0.21)	(0.19)	(0.01)	(0.20)	(0.16)	(0.01)
Heuristic 3 (Pay Down Highest Balance)	34.23	28.80	0.31	34.81	23.02	0.25
	(0.17)	(0.16)	(0.01)	(0.23)	(0.19)	(0.01)
Heuristic 4 (Pay Down Lowest Balance)	36.39	30.32	-0.10	30.20	19.03	0.28
	(0.20)	(0.18)	(0.01)	(0.21)	(0.17)	(0.01)
iii) Machine Learning Models						
Decision Tree	15.58	11.62	0.49	19.94	14.92	0.57
	(0.11)	(0.09)	(0.01)	(0.12)	(0.09)	(0.01)
Random Forest	13.47	9.71	0.66	16.79	11.25	0.73
	(0.11)	(0.08)	(0.01)	(0.12)	(0.09)	(0.00)
XGBoost	14.16	10.53	0.61	17.78	12.58	0.68
	(0.11)	(0.08)	(0.01)	(0.12)	(0.09)	(0.00)

Note: Goodness-of-fit for different models of the percentage of payments on the high-APR card. The first column shows the Root Mean Square Error (RMSE), the second column shows the Mean Absolute Error (MAE), and third column shows the Pearson Correlation Coefficient, which can also be interpreted as the square root of the R-squared. Round and non-round samples are defined by whether repayments on the high APR card are multiples £50. See Footnote 35 for details. The two-card sample is restricted to individual × months in which individuals face an economically meaningful allocative decision. Goodness of fit is calculated using the 20 percent hold-out sample and standard errors are constructed by the bootstrap method. See Section 2.2 for details on the sample construction.

Table A8: Heterogeneous Types from 3-Way and 4-Way Horse Race Models

Win Percent	(1)	(2)
Optimal	20.13	18.46
1/N	24.10	18.22
Balance Matching	55.77	49.10
Uniform		14.23

Note: Table shows percentage of individual × month observations that are best fit by different models of repayment behavior. The target variable is the share of repayments on the high APR card. All results shown in the table are based on the 20 percent hold-out sample. See Section 2.2 for more details on the sample construction.

Table A9: Correlation Matrix of Input Variables to Machine Learning Models

	APR(H)	APR(L)	Bal(H)	Bal(L)	Pur(H)	Pur(L)	Lim(H)	Lim(L)
APR(H)	1.00							
APR(L)	0.49	1.00						
Bal(H)	0.14	0.14	1.00					
Bal(L)	0.12	0.11	0.36	1.00				
Pur(H)	-0.05	-0.05	0.05	0.08	1.00			
Pur(L)	-0.05	-0.02	0.07	0.05	0.04	1.00		
Lim(H)	-0.01	0.04	0.61	0.23	0.16	0.08	1.00	
Lim(L)	-0.07	0.06	0.23	0.64	0.09	0.13	0.36	1.00

Note: Table shows correlation matrix for the input variables to the machine learning models. APR is the Annual Percentage Rate, Bal is the balance, Pur is purchases, and Lim is the credit limit. (H) indicates the high APR card and (L) indicates the low APR card. The two-card sample is restricted to individual $\times$ months in which individuals face an economically meaningful allocative decision. See Section 2.2 for details on the sample construction.

Table A10: Machine Learning Models Variable Importance

(1) Decision Tree		(2) Random Forest		(3) Extreme Gradient Boost	
Variable	Importance	Variable	Importance	Variable	Importance
Low Card Balance	0.21	High Card Balance	0.21	High Card Balances	0.25
High Card Balance	0.19	Low Card Balance	0.18	Low Card Balances	0.24
Low Card Credit Limit	0.13	High Card Credit Limit	0.13	High Card Purchases	0.19
High Card Credit Limit	0.12	Low Card Credit Limit	0.12	Low Card Purchases	0.17
Low Card Purchases	0.16	High Card Purchases	0.11	Low Card Credit Limit	0.06
High Card Purchases	0.18	Low Card Purchases	0.11	High Card Credit Limit	0.04
Low Card APR	0.00	High Card APR	0.07	Low Card APR	0.03
High Card APR	0.01	Low Card APR	0.07	High Card APR	0.02

Note: Table summarizes the importance of input variables in explaining payments on the high APR card in decision tree, random forest and extreme gradient boosting models. Rows show the proportion of the total reduction in sum of squared errors in the outcome variable resulting from the split of each variable across all nodes and all trees.

Table A11: Sensitivity Estimates Machine Learning Models Variable Importance

(1) Decision Tree			(2) Random Forest			(3) Extreme Gradient Boost		
Variable	Min	Max	Variable	Min	Max	Variable	Min	Max
Low Card Balance	0.18	0.26	High Card Balance	0.21	0.22	Low Card Balances	0.24	0.25
High Card Balance	0.15	0.19	Low Card Balance	0.20	0.20	High Card Balances	0.23	0.25
Low Card Credit Limit	0.12	0.18	Low Card Purchases	0.12	0.12	High Card Purchases	0.16	0.17
High Card Credit Limit	0.10	0.11	Low Card Credit Limit	0.11	0.12	Low Card Purchases	0.15	0.16
Low Card Purchases	0.09	0.18	High Card Purchases	0.11	0.12	Low Card Credit Limit	0.06	0.08
High Card Purchases	0.11	0.20	High Card Credit Limit	0.10	0.11	High Card Credit Limit	0.05	0.05
Low Card APR	0.00	0.03	High Card APR	0.07	0.07	Low Card APR	0.03	0.04
High Card APR	0.00	0.03	Low Card APR	0.06	0.07	High Card APR	0.03	0.03

Note: Table summarizes the importance of input variables in explaining payments on the high APR card in decision tree, random forest and extreme gradient boosting models. Rows show the proportion of the total reduction in sum of squared errors in the outcome variable resulting from the split of each variable across all nodes and all trees. The min and max values are the minima and maxima from machine learning models ran on 10 partitions of the 80 percent training sample used in Table A10.

Table A12: Interest Savings from Optimal Dynamic Model

Panel (A): Optimizing Payments and Spending

			Percentiles				
	Mean	Std. Dev.	10th	25th	50th	75th	90th
Savings in £							
3 Months	4.09	6.23	0.16	0.75	2.17	5.16	10.06
6 Months	15.25	20.03	0.66	2.97	8.64	20.53	38.20
9 Months	33.22	43.68	1.46	6.34	18.65	44.56	84.47
12 Months	57.78	76.53	2.55	10.80	32.01	77.39	148.74

Panel (B): Optimizing Payments

			Percentiles				
	Mean	Std. Dev.	10th	25th	50th	75th	90th
Savings in £							
3 Months	2.96	1.12	0.05	0.40	0.65	3.46	7.30
6 Months	12.00	16.35	0.45	1.61	7.38	18.81	24.21
9 Months	27.45	36.73	0.06	4.62	16.50	35.65	71.38
12 Months	38.45	50.26	0.83	7.59	21.93	50.88	99.28

Note: Table shows accumulated interest savings at 3, 6, 9 and 12 months from the optimal dynamic model. Savings are calculated as actual interest due minus interest due from the optimal dynamic model. Panel A shows savings from optimizing both payments and spending, Panel B shows savings from optimizing payments only. Two-card sample restricted to individual × months in which individuals face an economically meaningful allocative decision and then remain in the unrestricted data sample for 11 consecutive months. See Section 2 for details.

Table A13: Correlations Between Payment Rules

Panel (A): Balance Matching vs. Min. Pay Matching

	(1) Same Slopes	(2) Different Slopes	(3) Floor
Correlation	0.96	0.86	0.56
	(0.00)	(0.00)	(0.02)

Panel (B): Balance Matching vs. Actual

	(1) Same Slopes	(2) Different Slopes	(3) Floor
Correlation	0.63	0.41	0.50
	(0.00)	(0.00)	(0.00)

Panel (C): Min. Pay Matching vs. Actual

	(1) Same Slopes	(2) Different Slopes	(3) Floor
Correlation	0.61	0.28	0.23
	(0.00)	(0.01)	(0.02)

Note: Table shows correlation coefficients (standard errors in parenthesis) between balance-matching payments, minimum-payment-matching payments, and actual payments on the high APR. "Same Slopes" sample is account × months in which the minimum payment is determined by the percentage formula on both cards, and the percentage is identical across cards."Different Slopes" sample is account × months in which the minimum payment is determined by the percentage formula on both cards and the percentage differs across cards "Floor" sample is account × months in which the minimum payment determined by the floor value on both cards held by the individual, e.g. £25.

Table A14: Summary Statistics for Autopay (11 Percent of Observations) and Non-Autopay (77 Percent of Observations) Samples

	(1) Both Cards Non-Autopay	(2) Both Cards Autopay
i) Actual and Optimal Payments		
Actual Payments (Percent)	51.21	51.11
Optimal Payments (Percent)	74.36	58.42
Actual - Optimal Payments (Percent)	23.15	7.30
ii) Actual and Optimal Excess Payments		
Actual Excess Payments (Percent)	51.29	52.26
Optimal Excess Payments (Percent)	96.97	97.73
Actual Excess - Optimal Excess Payments (Percent)	45.68	45.47

Note: Table summarizes actual and optimal payments, and actual and optimal payments in excess of minimum due. The autopay sample is defined as observations where individuals make automatic payments on both cards. The non-autopay sample is defined as observations where individuals do not make automatic payments on either card. The two-card sample is restricted to individual × months in which individuals face an economically meaningful allocative decision. See Section 2.2 for details on the sample construction.

Table A15: Goodness-of-Fit for Different Models, Autopay and Non-Autopay Samples

	Both Cards Non-Autopay			Both Cards Autopay		
	(1) RMSE	(2) MAE	(3) Corr	(4) RMSE	(5) MAE	(6) Corr
i) Main Models						
Uniform Draw (0,100)	34.04	28.36	-0.01	36.90	30.30	-0.00
	(0.19)	(0.19)	(0.01)	(0.11)	(0.10)	(0.00)
Optimal	36.40	30.65	0.25	32.86	20.81	0.35
	(0.18)	(0.17)	(0.01)	(0.24)	(0.20)	(0.01)
1/N	17.64	12.63	-0.01	22.99	18.60	-0.00
	(0.10)	(0.10)	(0.01)	(0.06)	(0.07)	(0.00)
Balance Matching	22.00	16.81	0.38	23.11	15.61	0.53
	(0.14)	(0.11)	(0.01)	(0.13)	(0.09)	(0.00)

Note: Goodness-of-fit for different models of the percentage of payments on the high APR card. The first column shows the Root Mean Square Error (RMSE), the second column shows the Mean Absolute Error (MAE), and third column shows the Pearson Correlation Coefficient, which can also be interpreted as the square root of the R-squared. The autopay sample (11 percent of observations) is defined as observations where individuals make automatic payments on both cards. The non-autopay sample (77 percent of observations) is defined as observations where individuals do not make automatic payments on either card. The two-card sample is restricted to individual $\times$ months in which individuals face an economically meaningful allocative decision. Goodness of fit is calculated using the 20 percent hold-out sample and standard errors are constructed by the bootstrap method. See Section 2.2 for details on the sample construction.

Table A16: Horse Races Between Alternative Models, Autopay and Non-Autopay Samples

Panel (A): Uniform vs. Balance Matching

	Both Cards Non-Autopay (1)	Both Cards Autopay (2)
Win Percent		
Uniform	32.46	21.29
Balance Matching	67.54	78.71

Panel (B): Balance Matching vs. Optimal

	Both Cards Non-Autopay (1)	Both Cards Autopay (2)
Win Percent		
Balance Matching	75.21	61.02
Optimal	24.79	38.98

Note: Table shows percentage of individual × month observations that are best fit by different models of repayment behavior. The target variable is the share of repayments on the high APR card. Panel A compares balance-matching repayments against the lower benchmark where the percentage of repayments on the high APR card is randomly drawn from a uniform distribution with support on the 0 percent to 100 percent interval. Panel B compares optimal model repayments to the balance-matching model. The autopay sample (11 percent of observations) is defined as observations where individuals make automatic payments on both cards. The non-autopay sample (77 percent of observations) is defined as observations where individuals do not make automatic payments on either card. Samples are restricted to individual × months in which individuals face an economically meaningful allocative decision. All results shown in the table are based on the 20 percent hold-out sample. See Section 2.2 for more details on the sample construction.

References

Breiman, Leo. 2001. "Random Forests." *Machine Learning*, 45(1): 5–32.

Friedman, Jerome H. 2001. "Greedy function aproximation: A gradient boosting machine." *The Annals of Statistics*, 29(5): 1189–1232.

Friedman, Jerome, Trevor Hastie, and Robert Tibshirani. 2001. *The Elements of Statistical Learning*. Vol. 1, Springer Series in Statistics New York.

Keys, Benjamin J., and Jialan Wang. 2018. "Minimum Payments and Debt Paydown in Consumer Credit Cards." *Journal of Financial Economics*, (forthcoming).