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

John Gathergood† Neale Mahoney‡ Neil Stewart§ Jörg Weber¶

December 20, 2017

Abstract

We study how individuals repay their debt using linked data on multiple credit cards from five major issuers. We find that individuals do not allocate repayments to the higher interest rate card, which would minimize the cost of borrowing. Instead, individuals allocate repayments using a balance-matching heuristic under which the share of repayments on each card is matched to the share of balances on each card. We show that balance matching captures more than half of the predictable variation in repayments, performs substantially better than other models, and is highly persistent within individuals over time. Consistent with these findings, we show that machine learning algorithms attribute the greatest variable importance to balances and the least variable importance to interest rates in predicting repayment behavior.

Keywords: credit cards, consumer borrowing, rational behavior, balance matching, heuristics

JEL Codes: D12, D14, G02, G20

* We thank Joanne Hsu, Ben Keys, David Laibson, Brigitte Madrian, Devin Pope, Hiro Sakaguchi, Abby Sussman, Dick Thaler and seminar participants at Berkeley Haas, CEPR European Conference on Household Finance, Chicago Booth, FDIC, St Gallen, Luxembourg School of Finance, NBER Household Finance, Nottingham, Olso Applied Microeconomics Conference, and the Philadelphia Fed for helpful comments. Hanbin Yang provided excellent research assistance. This work was supported by Economic and Social Research Council grants ES/K002201/1, ES/N018192/1, ES/P008976/1, and Leverhulme grant RP2012-V-022.

† University of Nottingham, School of Economics; Network for Integrated Behavioural Science. Email: john.gathergood@nottingham.ac.uk.

‡ University of Chicago, Booth School of Business and NBER. Email: neale.mahoney@gmail.com.

§ University of Warwick, Department of Psychology. Email: neil.stewart@wbs.ac.uk.

¶ University of Nottingham, School of Economics; Network for Integrated Behavioural Science; Centre for Decision Decision Research and Experimental Economics. Email: joerg.weber@nottingham.ac.uk.

1 Introduction

Borrowing decisions underpin a broad set of economic behavior. Individuals borrow to smooth their consumption over the life-cycle, invest in human capital, and purchase durable goods, among other reasons. Thus, understanding how individuals borrow is (i) an important input for many fields of economic research and (ii) directly relevant for consumer financial policy.

This paper aims to shed light on this question by studying how individuals choose to repay debt – and thus implicitly how to borrow – across their portfolio of credit cards. We have a dataset with rich information on credit card contract terms, monthly statements, and repayments for 1.4 million individuals in the United Kingdom over a two-year period. Unlike other leading credit card datasets, our data allows us to link together multiple credit card accounts held by the same individual.[1] We study how individuals choose to allocate repayments across their credit cards.[2]

The credit card repayment decision is an ideal laboratory for studying borrowing because behavior that minimizes interest charges – what we refer to as optimal behavior – can be clearly defined. Consider individuals with debt on exactly two cards. Holding the total amount repaid on both cards in a particular month fixed, it is optimal for these individuals to make the minimum payment on both cards, repay as much as possible on the high interest rate card, and only allocate further payments to the low interest rate card if they are able to pay off the high interest rate card in full. What sets the credit card repayment decision apart from many other financial decisions is that optimal behavior does not depend on preferences, such as risk preferences or time preferences.[3] This allows us to evaluate models of behavior without having

[1] For instance, neither the OCC's Consumer Credit Panel nor the CFPB's Credit Card Database are designed to permit linking of accounts held by the same individual. The credit bureau datasets that combine information from multiple accounts held by the same individual do not have information on interest rates or repayments. There are a number of opt-in panels such as the Mint.com data and Lightspeed Research's "Ultimate Consumer Panel" that have information on multiple cards, but only for a self-selected sample of individuals.

[2] This type of allocative decision is common. In the U.K. market that we study, 46.1% of credit card holders have two or more cards, and this group accounts for 72.2% of outstanding balances (FCA, 2016). In the U.S. market, 71.5% of credit card holders have two or more cards, and this group accounts for 91.8% of balances (authors' calculations using a representative sample of 2015 TransUnion credit bureau data).

[3] For example, optimal mortgage choices depend on risk preferences (in the decision to use an adjustable or fixed rate mortgage) and time preferences over the real option to refinance in the future (see, Campbell and Cocco, 2003). There are very few institutional settings in which optimal mortgage choices can be clearly defined, such as in the Danish mortgage market (see, Andersen et al., 2017). The optimal credit card *spending* allocation depends on rewards programs, such as cash-back or airline points. Even when the terms of the rewards program are known, the optimal spending allocation depends on individuals' (idiosyncratic) value of features.

to jointly estimate preference parameters.

We start by showing that Ponce et al.'s (2017) finding of non-optimal credit card borrowing in Mexico is highly robust to the U.K. credit card market where we have data. Our baseline analysis focuses on individuals who hold exactly two cards in our dataset. For these individuals, the average difference in Annual Percentage Rate (APR) between the high and low interest rate cards is 6.3 percentage points, approximately one-third of the average 19.7% APR in our sample. If these individuals were completely unresponsive to interest rates, it is natural to assume that they would allocate 50.0% of their payments to each card on average. To minimize interest charges, we calculate that individuals should allocate 97.1% of the payments in excess of the minimum to the high APR card.[4] We show that individuals allocate only 51.5% of their excess payments to the high APR card, behavior that is virtually indistinguishable from the completely non-responsive baseline. In other words, 85% of individuals should put 100% of their excess payments on the high interest rate card but only 10% do so. Establishing this result is not the main focus of our analysis, but a necessary first step before going on to investigate alternative models.[5]

If individuals do not optimally allocate their credit card repayments, what explains their repayment behavior? One potential explanation is that individuals face a fixed cost of optimization – such as the time, psychological, or cognitive costs associated with determining the optimal repayment allocation (Sims, 2003). For some individuals, the reduction in interest costs may be too low to rationalize this fixed cost of optimization. We show, however, that the share of misallocated repayments is invariant to the difference in interest rates across cards (which can be as large as 15 percentage points) and to the size of the repayment amount (which can be as high as £800 in a month). The observed behavior, thus, seems inconsistent with a fixed-cost model of optimization.[6]

The main contribution of this paper is to evaluate heuristics that might better explain the

[4] The number is not exactly 100% because sometimes individuals can pay off the full balance by allocating a smaller amount, in which case they should allocate the remaining amount to the low interest rate card.

[5] In Section 3, we show that this result extends to the samples where we observe individuals allocating repayments across 3, 4, and 5 cards.

[6] The degree of misallocation is invariant to the time since account opening, indicating that learning cannot explain the observed behavior. Misallocation does not vary with the number of days between payment dates, indicating that frictions in co-ordinating repayments across cards cannot explain the observed behavior.

observed allocation of credit card payments. We first consider a balance-matching heuristic under which individuals match the share of repayments on each card to the share of balances on each card. The balance-matching heuristic naturally arises from the salient display of balances on credit card statements and the broad tendency for humans (and other species) to use "matching" heuristics in decision-making (discussed below). We also consider four alternative heuristics, such as the "debt snowball method" (under which payments are concentrated on the card with the lowest balance), which is recommended by some financial advisors.

We assess the explanatory power of these different repayment models using standard measures of goodness-of-fit (root mean square error, mean absolute error, Pearson's ρ). To provide a lower benchmark, we calculate goodness-of-fit under the assumption that the percentage of repayments on the high APR card is randomly drawn from a uniform distribution with support on the 0% to 100% interval. To provide an upper benchmark, we use machine learning techniques to find the repayment model that maximizes out-of-sample fit using a rich set of explanatory variables.

We find that balance matching captures more than half of the "predictable variation" in repayment behavior. That is, based on the range determined by the lower benchmark of random repayments and the upper benchmark of the machine learning models, we find that balance matching is closer to the upper benchmark on all of our measures. We also show that the optimal repayment rule and the other heuristic models do not come close to balance matching in their ability to match the data, capturing less than a quarter of the predictable variation for most measures.

In addition to providing us with an upper benchmark, the machine learning models allow us to assess the relative importance of interest rates versus balances in predicting repayment behavior. Consistent with the poor fit of the optimal repayment rule, we find that interest rates have low variable importance (i.e., partial R-squared) in our machine learning models. Consistent with the balance matching results, we find that balances have the highest variable importance, with importance factors substantially larger than any of the other explanatory variables. Unlike some other machine learning applications (e.g., Mullianathan and Spiess, 2017), these results are robust across partitions of the data.

We also evaluate each of our models in "horse race" type analysis where we determine the best fit model on an individual × month basis. In binary tests, balance matching has the best fit for twice as many observations as either the random, optimal, or other heuristics models, and balance matching performs comparably to the machine learning models. We also show that balance matching exhibits a high degree of persistence within individuals over time, suggesting that balance matching is more than a good statistical model but is actually capturing a stable feature of individual decision-making.

Our findings are related to a number of strands of literature. Our result on non-optimal repayments is closely related to the aforementioned Ponce et al. (2017) study, which finds that borrowing is highly non-optimal using linked data from Mexico. The first, modest, contribution of our paper is to show that the non-optimal behavior documented in Ponce et al. (2017) in Mexico extends to the U.K. market that we study.[7]

Our main result on balance matching relates to a literature in psychology and economics on heuristics in individual decision-making. The fact that individuals focus on balances, which are prominently displayed on credit card statements, connects to a literature on how saliently placed information can provide an anchor for choices (e.g., Tversky and Kahneman, 1974; Ariely et al., 2003; Stewart, 2009; Bergman et al., 2010).

Our finding on balance matching also shares a resemblance with a long line of research on probability matching. Herrnstein's (1961) *matching law* is based on the observation that pigeons peck keys for food in proportion to the time it takes the keys to rearm rather than concentrating their effort on the key that rearms most quickly. Rubinstein (2002) shows, in an experimental study, that subjects diversify when choosing between gambles with a 60% and 40% chance of winning, even though the option with a 60% chance of winning dominates any other strategy (see Vulkan, 2000 for a review of this literature). Balance matching is also reminiscent of the classic Benartzi and Thaler (2001) result on how investors in defined-contribution saving plans allocate funds such that the proportion invested in stocks depends strongly on the proportion of stock funds in the choice set.[8]

[7] Our paper is more broadly related to a large literature on price dispersion and non-optimal behavior in consumer financial markets. See, for example, Keys et al. (2016), Hortaçsu and Syverson (2004), and Stango and Zinman (2016).

[8] See DellaVigna (2009) for a review of the evidence on choice heuristics using field data.

The caveats to our analysis largely stem from the fact that we focus on the allocative decision of how individuals split repayments across their portfolio of credit cards. While this decision greatly simplifies the analysis, our estimates of the degree of non-optimal behavior should be interpreted as lower bounds relative to a counterfactual in which individuals could additionally reallocate payments across non-credit card loans (such as mortgages or automobile loans) or make adjustments on the extensive margin (e.g., by adjusting the tradeoff between debt repayment and consumption). Our focus on the allocative decision also naturally leads us to consider "allocative heuristics," such as balance matching, rather than heuristics that determine behavior on the extensive margin. For example, balance matching could arise from individuals repaying a fixed percentage of their balances (e.g., 10%), a rule-of-thumb that would lead to inefficient behavior on both the allocative and extensive margins.

Finally, it is important to point out that we do not view balance matching as a "model" of behavior in the psychological or structural-economic sense of the term. Instead, we think about balance matching as a simple allocation rule which provides a useful statistical approximation of behavior. Balance matching is easy to understand, reinforces existing theories of behavior (e.g., probability matching, Herrnstein's matching law), and might provide intuition in yet-to-be-studied environments. Understanding the psychological underpinnings of balance matching is a question we leave for future work.

The rest of the paper proceeds as follows. Section 2 describes our data and presents summary statistics for our baseline sample. Section 3 presents our results on the optimality of repayment behavior. Section 4 examines rounding and a $1/N$ rule for repayments. Section 5 lays out alternative heuristics for debt repayment, including the balance-matching heuristic. Section 6 tests between these repayment models. Section 7 presents sensitivity analysis. Section 8 concludes.

2 Data

2.1 Argus Credit Card Data

Our data source is the Argus Information and Advisory Services' "Credit Card Payments Study" (CCPS). The CCPS has detailed information on contract terms and billing records from five major credit card issuers in the U.K. These issuers have a combined market share of over 40% and represent a broad range of credit card products and market segments. We have obtained monthly data covering January 2013 to December 2014 for a 10% representative sample of individuals in the CCPS who held a credit card with at least one of the five issuers. Unlike other leading credit card datasets, the CCPS provides us with anonymized individual-level identifiers that allow us to link together multiple accounts held by the same individual.[9]

2.2 Sample Restrictions

Our interest lies in understanding how individuals make repayment decisions across their portfolio of credit cards. Holding multiple cards is not uncommon. In the U.K. market, 46.1% of credit card holders have two or more cards, and individuals with two or more cards account for 72.2% of outstanding balances (FCA, 2016).[10]

Our unit of analysis is the individual × month. In the remainder of the paper, we refer to individual × months interchangeably as "observations". All of the credit cards in our data require payments at a monthly frequency. We consider cards to be in the same "month" if their billing cycles conclude in the same calendar month. Since billing cycles often conclude near the end of the calendar month, payment dates are often quite near to each other.[11] We construct separate samples based on the number of credit cards held by the individual in that month in our dataset (e.g., two cards, three cards, and so on). This is a weak lower bound on the number of cards held by the individual. In the analysis that follows, we typically start by presenting our methodology and results for the two-card sample, and then examine how our findings extend to individuals with three or more cards.

To focus on individual repayment decisions, we first clean the sample by dropping observations where individuals are delinquent or have defaulted on at least one card in their

[9] As discussed in Footnote 1, neither the OCC's Consumer Credit Panel nor the CFPB's Credit Card Database are designed to permit matching of multiple individually-held accounts, and credit bureau datasets typically do not have information on interest rates or repayments. Opt-in panels such as Lightspeed Research's "Ultimate Consumer Panel" have information on multiple cards, but only for a self-selected sample of individuals.

[10] These numbers are even higher in the U.S. market. Using a representative sample of 2015 TransUnion data, we calculate that 71.5% of credit cards holders have two or more cards, and individuals with two or more cards account for 91.8% of balances and 91.7% of revolving balances.

[11] In the two-card sample, two-thirds of observations have payment dates that are 10 days apart or fewer. See Figure A5 for a histogram of the difference in due dates between cards.

portfolio, or where individuals pay less than the minimum due or more than the full balance on at least one card.[12] Together these restrictions drop 2.0% of individuals and 4.2% of aggregate revolving debt from the two-card sample. Second, we focus on observations in which individuals hold debt on all of their cards – i.e., they are carrying "revolving" balances. This ensures that allocating repayments towards the high APR card, in the manner described in Section 3, minimizes interest charges.[13,14] These restrictions remove a further 25.0% of individuals from the two-card sample. However, since most of these individuals do not have revolving debt – i.e., they are "transactors" who repay the balances in full each month – we only drop 12.7% of aggregate revolving debt.

Third, we make a number of sample restrictions so that we focus on observations where individuals, holding fixed total monthly repayments, have scope to reallocate payments across cards and therefore face an *economically meaningful* allocative decision. In particular, we drop a small number of observations where the interest rate is identical across cards, since any reallocation of payments has no impact on the cost of borrowing. We then drop observations where the individual pays either the full balance or the minimum payment on all of their cards, since these individuals do not have any payments to reallocate. Taken together, this third set of restrictions drops an additional 35.4% of individuals and 24.9% of aggregate revolving debt in the two-card sample. Appendix Table A1 goes through these restrictions one-by-one. Most of the reduction is due to dropping individuals who pay exactly the minimum on each card.

Table 1 presents summary statistics on the baseline two-card sample. The average difference

[12] Paying less than the minimum or more than full balance sometimes results from "mistakes" that are difficult to interpret with an allocative model of behavior. For instance, zero payments sometimes result from "forgetting" to put a check in the mail. Similarly, overpayments sometimes result from refunds that are processed after the individual decides on their payment, which reduce the balance below the payment amount.

[13] A (complicated) feature of credit cards is that if an individual carries no revolving balance at the beginning of the month, and repays the balance in full, they avoid any interest charges that month. If an individual carries a revolving balance at the beginning of the month, interest charges are incurred on average daily balance irrespective of whether the card is repaid in full. We focus on individuals who begin the month with revolving balances on all cards as it is unambiguously interest-cost-minimizing for these individuals to allocate repayments towards the high APR card. In other scenarios, it could be interest-cost-minimizing to allocate repayments towards the low APR (although our back-of-the-envelope calculations suggest this is unlikely).

[14] One consequence of this restriction is that we omit individuals who have two cards but hold revolving balances on only one card. In doing so, one potential concern is that we drop individuals who have "fully optimized" by completely paying off their high interest rate card. If this were the case, then our sample would be selected on individuals who failed to optimize, raising issues of external validity. However, among individuals who carry debt on only one card, the majority (61.8%) carry debt on only the high interest rate card, indicating that our sample is not selected in this manner.

in APR (for purchases) between the high and low interest rate cards is 6.3 percentage points, or approximately one-third of the 19.7% average purchase APR in the sample.[15] Yet despite this substantial difference in prices, utilization is remarkably similar. Purchases are £128 on the high APR card versus £117 on the low APR card; repayments are £260 on the high APR card versus £230 on the low APR card; and revolving balances are £2,200 on the high APR card versus £2,054 on the low APR card. This is particularly striking given that credit limits are almost three times larger than revolving balances on average, indicating that the typical individual would be able to shift all of their borrowing to the low APR card without exceeding their credit limit.

3 Optimal Repayments

We start by comparing the actual and interest-cost-minimizing allocation of repayments across cards. We refer to the interest-cost-minimizing allocation as the "optimal" allocation because it is hard to think of a (reasonable) scenario where minimizing interest costs would not be optimal. Holding the total repayment amount on all cards fixed, it is optimal for individuals to make the 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.[16]

We focus on repayments, rather than other measures of credit card use like spending or revolving balances, because, for repayments, we can clearly define optimal behavior. In contrast, optimal spending may depend upon rewards programs, which we do not observe in our data.[17] We also do not focus on the optimality of revolving balance allocations because revolving balances are a "stock" that cannot typically be quickly adjusted.[18] Thus, to determine whether

[15] This difference does reflect short-term 0% promotional interest rate offers, which account for less than 5% of account × month observations in the baseline sample.

[16] We explicitly rule out the possibility that choosing not to make the minimum payment on a lower interest rate card could be optimal. Failing to repay the minimum repayment would result in a penalty fee and a marker on the individual's credit file.

[17] While issuers typically incur only a small cost for the rewards they provide – approximately 1%, see Agarwal et al. (2015) – individuals might value rewards (such as airline points) at a high enough value to affect optimal spending decisions.

[18] In particular, with the exception of balance transfer products in the prime credit card market, individuals can only reallocate their stock of revolving balances by adjusting the flow of spending and repayments on a month-by-month basis.

revolving balances are "optimal", we would need to take a stand on how individuals could reallocate revolving balances through counterfactual spending and repayment decisions over time, which would require us to know the individuals' time preference and their expectations over future spending and repayment decisions.

Panel A of Figure 1 plots the distribution of actual and optimal payments in the baseline two-card sample. The distribution of actual repayments appears close to symmetric, with a mass point at 50%, and smaller mass points at 33% and 67%. In contrast, the distribution of optimal repayments is heavily weighted towards the high APR card. It is not optimal for individuals to place 100% of their payments on the high interest rate card because (ii) they need to pay the minimum on the low interest rate card and (ii) they are sometimes able to pay off more than the full balance on the high interest rate card.

Summary data for actual and optimal repayments for the two-card sample is shown in Table 2. On average, individuals should allocate 70.7% of repayments to the high APR card. If individuals were completely unresponsive to interest rates, we might expect them to place 50% of payments on the high interest rate card. On average, individuals allocate 51.2% to the high interest rate card, which is very close to the completely non-responsive baseline. Individuals, thus, misallocate 19.5% of their total monthly payment on average.

In Figure A1 we plot misallocated repayments in *excess* of the minimum payment. That is, we subtract out the amount required to make the minimum payment on each card and then calculate the share of the remaining amount that is allocated across cards. On average, individuals should allocate 97.1% of payments in excess of the minimum to the high APR card, whereas in practice they actually allocate 51.5% to that card.[19] Alternatively put, 85% of individuals should put 100% of their excess payments on the high interest rate card but only 10% do so. Summary data for payments in excess of minimum are shown in Table A2.

Panels B to D of Figure 1 show radar plots of the average percentage of actual and optimal payments on each card for the samples with 3, 4, and 5 cards. In each of the plots, the cards are ordered clockwise from the highest to the lowest APR (starting at the first node clockwise from noon). The polygons for actual payments are symmetric, indicating that the actual percentage

[19] The number is not exactly 100% because sometimes individuals can pay off the full balance by allocating a smaller amount, in which case they should allocate the remaining amount to the low interest rate card.

of payments is very similar across cards. The polygons for optimal payments show that it would be optimal to allocate a substantially higher percentage of payments to the highest APR card and a substantially lower percentage to the card with the lowest APR.

3.1 Fixed Costs of Optimization

One potential explanation for the non-optimality of repayments is that individuals face a fixed cost of optimization (Sims, 2003). Specifically, in our context, individuals are already making positive credit card repayments on both cards and therefore already paying the fixed cost of logging into their bank's website or sending a check. So by a fixed cost of optimization, we having in mind the time, psychological, or cognitive costs associated with making the optimal repayment relative to making a non-optimal payment amount.[20] For some individuals, the reduction in interest payments from cost-minimizing may be too low to rationalize incurring this fixed cost.

To investigate this potential explanation, we examine the correlation between the percentage of misallocated repayments and the economic stakes of the repayment decision in the two-card sample. We define misallocated payments as the difference between optimal and actual payments on the high APR card. We examine two measures of the economic stakes: (i) the difference in APR across cards and (ii) the total repayments made that month. Since the gains from optimizing are increasing in the financial stakes, under the fixed cost explanation, the percentage of misallocated repayments should be declining in both measures. Moreover, for individuals with large economics stakes, we would expect the degree of misallocation to be close to zero.

Panel A of Figure 2 shows a binned-scatter plot of the percentage of misallocated payments against the difference in APR between the high and low interest rate cards. The binned-scatter plot is constructed by partitioning the x-axis variable into 20 equal-sized groups and plotting the mean of the y-axis and x-axis variables for each group.[21] The flat relationship indicates that individuals are not less likely to misallocate repayments even when there is a large APR

[20] This contrasts with optimal mortgage refinancing (e.g., Keys et al., 2016), a setting where optimization would require actively soliciting a new mortgage offer.

[21] See Chetty et al. (2014) for more details on the binned-scatter plot methodology.

difference (more than 15 percentage points).[22]

Panel B of Figure 2 shows a binned-scatter plot of the percentage of misallocated payments against total repayments on both cards. Again, there is no evidence of a decreasing relationship. Indeed, the relationship is increasing because individuals who make the largest payments can cover the minimum on the low APR card with a smaller percentage of their overall allocation and thus *should* allocate an even larger fraction of payments to the high APR card.[23]

An alternative potential explanation for the observed non-optimal behavior is that individuals learn over time (e.g., since opening a card), and that our analysis of the cross-sectional distribution of repayments masks this learning behavior. A model with time-varying adjustment costs (in the spirit of Calvo, 1983) would also generate a gradual reduction in the degree of misallocation over time. Panel C of Figure 2 examines this explanation by showing a binned-scatter plot of the percentage of misallocated payments against the age (in months) of the high APR card. For this analysis, we restrict the sample to individuals who open a high APR during our sample period and for whom we can observe economically meaningful allocation decisions for 10 consecutive months. In the plot, the horizontal axis starts in the second month after opening, since this is the first month in which individuals could have a balance on the high APR card to repay. The plot shows no evidence of a reduction in the percentage of misallocated repayments over time. This finding suggests that neither learning nor time-varying adjustment costs can explain the observed non-optimizing behavior.[24]

A somewhat different potential explanation for the observed behavior is that individuals face within-month liquidity constraints that prevent them from optimally reallocating payments. For example, individuals might receive their paycheck between credit card due dates, and therefore have different amounts of cash-on-hand when making their payments.[25] Panel D of Figure 2 examines this explanation by showing a binned-scatter plot of the percentage

[22] Panel A of Figure A1 shows that there is a similarly flat relationship between misallocated payments in excess of the minimum payment and the difference in APR across cards.

[23] Panel B of Figure A1 illustrates the relationship between misallocated payments in excess of the minimum payment and the total repayment across both cards. There is a slight downward slope, but certainly not the type of relationship that would be predicted by a fixed-cost-of-optimization model.

[24] Panel C of Figure A1 illustrates the relationship between misallocated payments in excess of the minimum payment and the age (in months) of the high APR card.

[25] For instance, if the due date on the high APR card occurred before the individual received their paycheck and the due date on the low APR card occurred after, it might be optimal to repay a larger amount of the lower APR card balance.

of misallocated payments against the difference in payment due dates. The flat relationship indicates that this type of within-month allocation friction cannot explain our finding.[26]

3.2 Costs of Misallocation

What are the costs of the failure to optimize? We consider 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. For the two card sample, the optimal allocation can be achieved by transferring balances from the high to the low interest rate card up to the point where the individual "maxes out" their low interest rate card. With multiple cards, the optimal strategy is to allocate as much of the aggregate balance as possible to the credit card with the lowest interest rate, then allocate any remaining balance to the card with the next lowest rate, and so on.

Table 3 presents the annualized interest savings from this exercise. Average interest savings are increasing across the number of cards, rising from £65 in the two-card sample to £248 in the five-card sample. Because the degree of misallocation is not declining in the economic stakes of the decision, individuals with larger balances and larger differences in interest rates have a substantial cost of misallocation, with the 90th percentile rising from £167 in the two-card sample to £927 in the five-card sample.

There are a number of caveats to this analysis. Since our sample covers approximately 40% of the market, we likely only observe a subset of an individuals' credit card portfolio. Allowing individuals to optimize over more cards would necessarily lead to a larger cost of misallocation. And of course, allowing individuals to optimize across different types of debt or across savings and borrowing products would lead to even larger values. The interest savings we calculate are larger than the single period savings from optimal repayments because they are annualized (multiplied by 12) and because (absent a balance transfer) the counterfactual steady state can only be achieved through optimal repayments and spending decisions over a number of periods.

[26] Panel D of Figure A1 shows that there is similarly no relationship between misallocated payments in excess of the minimum payment and the difference in due dates

4 Rounding and the $1/N$ Rule

The spike in repayments at 50% (see Panel A of Figure 1) suggests that some individuals use a simple $1/N$ heuristic in which they make equal-sized repayments across cards, analogous to the $1/N$ heuristic documented in defined-contribution savings decisions (Benartzi and Thaler, 2001). In particular, the excess mass of individuals who make payments at a one-to-one ratio is approximately 8.2%. If we add in the excess mass at one-to-two and two-to-one ratios, we can explain 11.7% of repayments.

While 11.7% is fairly modest, it is not clear whether even this percentage of individuals is behaving according to an $1/N$-type model. As we show below, a fairly sizable fraction of individuals round payments to £50, £100, £200, and so on. If an individual rounds up a payment on card A from £80 to £100 and rounds down a payment on card B from £120 to £100, then the individual would appear as if they intended to make equal-sized payments, even though, absent rounding, the share of payments on each card would be substantially different from 50%.

Figure 3 investigates this competing explanation for the spike at 50%. Panel A plots the distribution of payments on the high APR card in pounds, and shows substantial evidence of rounding. We calculate that 19.2% of payments take on values that are multiples of £100, and 33% of payments take on values that are multiples of £50 (which obviously includes payments that are multiples of £100). Panels B and C show the percentage of payments on the high APR card, splitting the sample by whether the individual makes round number payments (defined as multiples of £50) or "non-round" number payments on the high APR card.[27]

The plots show that the peaks at 50% (as well as 33% and 66%) are heavily concentrated among individuals who make round number repayments. In the non-round sample, there is only a small spike at 50%, and no discernible spike at 33% or 66%. In other words, virtually nobody chooses to make equal payments of any other value that is not a multiple of £50. If individuals actually wanted to repay equal amounts, it is hard to imagine that we would not

[27] The propensity to make round number payments is highly correlated across cards within an individual × month. We calculate that 73% of individuals who make a round number payment on the high APR card also make a round number payment on the low APR card, and 78% of individuals who make a non-round number payment on the high APR card also make a non-round number payment on the low APR card. Dividing the sample into individuals who make round number repayments on both cards, non-round number repayments on both cards, and a mix of round and non-round repayments complicates the exposition without changing the results.

see this type of repayment behavior.

We therefore do not view the spike at 50% as compelling evidence of a $1/N$-type heuristic. Indeed, we cannot reject the view that nearly all of the spike at 50% is due to rounding. However, since we do not have random variation in whether individuals round, we cannot rule out the possibility that some of these individuals would have allocated 50% on the high APR card if they had counterfactually not rounded their payments. Thus, we conclude that the percentage of individuals who behave according to a $1/N$-type rule is bounded between 0% and 11.7%.

5 Balance Matching and Other Heuristics

If individuals do not optimally allocate their repayments (and a $1/N$-type rule can explain only a small fraction of behavior), what explains credit card repayments? In the remainder of this paper, we evaluate heuristics that might better explain the allocation of repayments. In this section, we introduce the set of heuristics that we consider. In Section 6, we evaluate the explanatory power of these heuristics.

5.1 Balance Matching

We first consider a balance-matching heuristic under which individuals match the share of repayments on each card to the share of balances on each card. Let $k = 1 \ldots K$ index cards, p_k indicate payments, and q_k indicate balances. Balance-matching payments are given by

$$\frac{p_k}{\sum_{\kappa=1}^{K} p_\kappa} = \frac{q_k}{\sum_{\kappa=1}^{K} q_\kappa} \qquad \text{for} \quad k = 1 \ldots K \tag{1}$$

subject to the constraint that the individual pays at least the minimum and no more than the full balance on any of their cards.[28]

The balance-matching heuristic naturally arises from the salient display of balances on credit card statements and the broad tendency for humans (and other species) to adopt "matching" heuristics in decision-making. As shown in Figure 4, balances are perhaps the most

[28] In the two-card sample, only 13.0% of observations are affected by these constraints. In nearly all of these cases, the balance-matching payment is less than 2 percentage points below the minimum payment amount. Treating these observations in other ways (e.g., dropping these observations) does not have a material impact on the results.

prominently displayed element on credit card statements. The psychological theory of *anchoring* (Tversky and Kahneman, 1974) suggests that individuals might make payments in relation to this saliently displayed balances (instead of less saliently displayed interest rates).[29,30]

A proclivity for "matching" has been observed across species and across domains, and thus may result from a deeper underlying tendency for proportionality in decision-making. Herrnstein's (1961) *matching law* is based on the observation that pigeons peck keys for food in proportion to the time it takes the keys to rearm rather than concentrating their effort on the key that rearms most quickly. The *probability matching* literature finds that individuals place bets in proportion to the probability of payoffs, even though betting on the option with the highest probability of payoff first-order stochastically dominants any other decision. For example, Rubinstein (2002) shows in an experimental study that subjects diversify across independent 60%-40% gambles even though betting on the gamble with a 60% probability of payout is a strictly dominant strategy. [31]

Of course, we do not propose balance matching as a precise description of individual repayment behavior. Pigeons do not measure the time it takes the keys to rearm with a stopwatch and we do not mean to suggest that individuals use long division to calculate the share of repayments that should be allocated to each card. Instead, we propose that individuals approximate balance matching in their repayment behavior. Indeed, since credit card balances are fairly stable over time, an individual could approximate a balance matching rule without knowing the exact balance on each card that month.

5.2 Other Heuristic Models of Repayment

We also consider four alternative heuristics that capture intuitive economic and non-economic approaches to the allocation of payments. Some of these heuristics are based on the *capacity* of a credit card, which we define as the difference between the credit limit and current balance

[29] A second reason why balances may be more salient is that balances are denoted in the same units as repayments (£s), whereas APR take on different units (%).

[30] Balances also enter the minimum payment formula. Therefore, at least in principle, repayments might depend on balances indirectly through the minimum payment amount. We discuss this issue in Section 7 and show that this channel does not explain our results.

[31] See Vulkan (2000) for a review of the matching literature and DellaVigna (2009) for a review of the evidence on choice heuristics using field data.

in pounds. We describe these heuristics for the two-card sample, but they could be naturally extended to settings with three or more cards.

- **Heuristic 1**: *Repay the card with the lowest capacity.* Allocate payments to the lowest capacity card, subject to paying the minimum on both cards. Once capacity is equalized across cards, allocate additional payments to both cards equally. Intuitively, by focusing payments on the card with the lowest capacity, this heuristic reduces the risk that an accidental purchase will put an individual over their credit limit, which would incur an over-limit fee and marker on the individual's credit file.
- **Heuristic 2**: *Repay the card with highest capacity.* Allocate payments to the highest capacity card, subject to paying the minimum on both cards. Once the highest capacity card is fully repaid, allocate remaining payments to the other card. Intuitively, by allocating payments to the card with the highest capacity, this heuristic creates maximum "space" for making a large purchase on a single card (e.g., buying a television).
- **Heuristic 3**: *Repay the card with the highest balance.* Allocate payments to the highest balance card, subject to paying the minimum on the other card. Once balances are equalized across cards, allocate additional payments to both cards equally. If individuals dislike having a credit card with a large balance, this heuristic reduces the maximum balance they are carrying, and thus might explain repayment behavior.
- **Heuristic 4**: *Repay the card with the lowest balance ("debt snowball method").* Allocate payments to the lowest balance card, subject to paying the minimum on the other card. Once the balance on the lowest balance card is paid down to zero, allocate any additional payments to the other card. This heuristic is sometimes referred to as the debt snowball method by financial advisors. Proponents argue that paying off a card with a low balance generates a "win" that motivates further repayment behavior. If an individual fully pays off a card, this heuristic has the additional benefit of "simplifying" the individual's debt portfolio.

6 Testing Repayment Models

We evaluate balance matching and the other heuristics using two statistical approaches. First, we measure the distance between observed and predicted behavior using standard measures of goodness-of-fit (root mean square error, mean absolute error, Pearson's ρ). Second, we evaluate the performance of our models in a "horse race" type analysis where we determine the best fit model on an individual $\times$ month basis.

6.1 Goodness-of-Fit

We start by presenting visual evidence on the goodness-of-fit of the balance-matching heuristic. Figure 5 examines fit in the two-card sample. The left column shows the marginal distributions of actual and balance-matching payments on the high APR card. The right column displays the joint distribution using a contour plot. The top row shows these relationships for the baseline sample, the middle row restricts attention to the sample with round number payments (defined as multiples of £50), and the bottom row focuses on the sample with non-round number payment amounts.

The histograms (left column) show that the marginal distributions of actual and balance-matching payments are quite similar, except for the spikes at 33%, 50%, and 66%. As we discussed in Section 4, the spikes are much smaller in the non-round sample, suggesting that these $1/N$-type repayments may be an unintended consequence of rounding, and not reflective of underlying repayment behavior. The higher mass along the 45 degree line in the joint densities (right column) indicates that actual and balance-matching payments are strongly correlated. The correlation is strongest in the non-round payment sample and weaker in the round number payment sample, with the horizontal streaks in the round number sample reflecting the excess mass at 33%, 50%, and 66%. The weaker correlation in the round payment sample is consistent with a two-stage model in which individuals first decide to make balance-matching payments and then add noise to the process by rounding their repayment amounts.

Figure 6 examines the fit of the balance-matching model in the samples with 3, 4, and 5 cards. The left column shows marginal distributions of the share of payments on the highest APR card. The right column shows radar plots of the average share of actual and balance-matching payments. In each of the radar plots, the cards are ordered clockwise from the highest to the lowest balance (starting at the first node clockwise from noon). The distribution of payments on the high APR card (left column) is similar to predicted payments under balance matching. The radar plots (right column) show that the average share of payments is fairly close to those predicted by balance matching, especially compared to the radar plots that examine optimal repayment behavior (Figure 1).

We formally measure the performance of the balance matching and alternative models using three standard measures of goodness-of-fit: the square root of the mean square error (RMSE), the mean absolute error (MAE), and the correlation between actual and predicted payments (Pearson's ρ).[32] To help interpret the goodness-of-fit values, we also establish lower and upper benchmarks. For a lower benchmark, we calculate goodness-of-fit under the assumption that the percentage of repayments on the high APR card is randomly drawn from a uniform distribution with support on the 0% to 100% interval. To provide an upper benchmark, we use machine learning techniques to construct a set of purely statistical models of repayment behavior. Specifically, we estimate decision tree, random forest, and extreme gradient boosting models of the percentage of payments allocated to the high APR card. We use the same set of variables which enter into our heuristics (APRs, balances, spending and credit limits on both cards) as input variables and "tune" the models to maximize out-of-sample power using standard methods from machine learning.[33] We use 80% of the data sample as the "training" sample and measure out-of-sample fit on a 20% "hold-out" sample. For consistency, in the analysis that follows, we compare all models using the hold-out sample. Technical details are provided in Appendix I.

Figure 7 shows these measures of goodness-of-fit under the different models. (Appendix Table A3 shows the underlying numerical amounts with bootstrapped standard errors, constructed by drawing with replacement from the hold-out sample.) The optimal model yields only a very small improvement in the RMSE and MAE relative to the lower benchmark. The

[32] Pearson's ρ is also the square root of the R-squared from a univariate regression of actual payments on predicted payments.

[33] We view these machine learning models as "prior-free" models of repayment behavior. If we additionally include our candidate models (optimal, balance matching, and the other heuristics) as input variables, we obtain only small improvements in model fit.

optimal model does generate an economically meaningful increase in the Pearson correlation, although this is partly because the lower benchmark (uniformly distributed random amount) is constructed to have a Pearson correlation of 0. The other heuristics perform similarly poorly, with the goodness-of-fit measures generally falling less than a quarter of the way between the lower and upper benchmarks.[34]

The goodness-of-fit measures for the balance-matching model fall slightly more than halfway between the lower and upper benchmarks, indicating that balance matching captures more than half of the "predictable variation" in repayment behavior. Appendix Figure A3 shows goodness-of-fit separately for the round and non-round samples (defined as multiples of £50). The balance-matching model captures a larger fraction of the predictable variation in the non-round number sample, which is consistent with a two-stage model of payments discussed above.

There are two ways to view the performance of the balance-matching model relative to the upper benchmark provided by the machine learning models. The glass half full view is that being able to capture more than half of the predictable variation in repayment behavior with a simple balance matching model is useful. Balance matching is a useful description of behavior because it is easy to understand, reinforces existing theories of behavior (e.g., probability matching, Herrnstein's matching law), and might provide intuition for individual behavior in yet-to-be-studied environments. The glass half empty perspective is that machine learning techniques provide higher predictive power. Thus, if the goal is prediction – rather than understanding human behavior – machine learning techniques may be preferable.

6.1.1 Balances and APRs in Machine Learning Models

In addition to providing us with an upper benchmark, the machine learning models allow us to asses the relative importance of balances and APRs in predicting repayment behavior.

[34] Appendix Figure A4 shows the marginal distributions of actual and predicted payments under each of the alternative heuristics. One common feature of these alternative heuristics is that they predict that individuals should often concentrate their excess payments on a single card. For instance, under Heuristic 1 (repay the card with the lowest capacity), individuals should fully allocate repayments, in excess of the minimum, to the card with the lowest capacity until the point where both cards have equal capacity remaining. Individuals, however, seem to avoid "corner solutions" in their repayment behavior. As a result, the alternative heuristics over-predicted the share of individuals who allocate a very small (less 10%) or very large (greater than 90%) share of payments to the high APR card.

Specifically we calculate the variable importance, which can be thought of as the incremental increase in R-squared from adding a given variable to the model. Appendix Table A5 shows that APRs and balances are not strongly correlated. In cases where the variables are collinear, the interpretation of variable importance may be spurious (Mullianathan and Spiess, 2017).

Consistent with the balance matching results, the machine learning models confirm that balances are hugely important for predicting behavior. Table 4 shows that balances have the highest variable importance in all of the models, with importance factors substantially larger than any of the other explanatory variables. Consistent with the poor fit of the optimal repayment rule, we find that APRs have the lowest variable importance across models. Appendix Table A6 shows minima and maxima of variable importance from models estimated on 10 partitions of the training sample. Unlike some other machine learning applications (e.g., Mullianathan and Spiess, 2017), the ranges are narrowly spread around the baseline estimates in Table 4, indicating that our variable importance measures are not particularly sensitive to random variation in the training data.

6.2 Horse Races Between Alternative Models

The goodness-of-fit analysis effectively measures the distance between observed repayments and predicted repayments using different loss functions. An alternative approach is to run "horse races" where we determine the best fit model on an observation-by-observation basis. A model that fits a small number of observations very poorly, but a larger number quite well, would perform poorly under most distance metrics (especially those with convex loss functions) but would perform well under this alternative approach.

6.2.1 Methodology

Let i denote individuals and t denote months. Let p_{it} indicate the actual share of payments that is allocated to the high APR card and let $\hat{p}^j_{it}$ indicate the share of payments on the high APR card predicted by model $j \in J$. To test between alternative models, we estimate specifications of the form:

$$p_{it} = \left(\sum_{j \in J} \lambda_{it}^{j} \, \hat{p}_{it}^{j} \right) + \epsilon_{it} \qquad \text{s.t.} \quad \lambda_{it}^{j} \in \{0, 1\} \quad \text{and} \quad \sum_{j} \lambda_{it}^{j} = 1, \tag{2}$$

where the λ_{it}^{j} are indicators that "turn on" for one and only one of the candidate models $j \in J$. We vary the set of alternatives J to allow for horse races among different competing sets of models. When there are multiple models that are tied for the closest to actual repayments, the λ_{it}^{j} are not identified, and so we drop these observations from our analysis. We estimate the model by minimizing the absolute deviation between the observed and predicted values.[35] Intuitively, for each observation p_{it}, this procedure picks the model j that best fits observed repayment behavior at the individual × month level.

It is worth pointing out that our ability to identify the best-fit model at the individual × month level is due to the unique nature of the credit card repayment decision. As discussed in Section 1, what sets credit card repayments apart from many other financial decisions is that optimal behavior does not depend on preferences (such as risk preferences or time preferences). If preferences were important for optimal behavior, then conducting this type of exercise would require recovering preferences at the individual level, which would be a significant empirical challenge.

6.2.2 Results

Table 5 shows results of this horse race analysis in the pooled sample of individual × months. Panel A compares each of our models one-by-one 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% to 100% interval. In a binary comparison, balance matching is the best fit model for 67.2% of observations, or about twice the percentage of the uniform benchmark. The optimal model and the other heuristics are the closest for slightly more than half of the observations, and therefore only perform slightly better than the uniform benchmark. In binary comparisons, the machine learning models have the best fit for between 69.3% and 76.6% of

[35] In practice, since we are estimating a separate set of coefficients for every individual × month, the estimates would be identical under a quadratic (or any other monotonically increasing) loss function.

observations, which is similar to balance matching.[36]

Panel B of Table 5 compares each of the models one-by-one to the balance-matching model.[37] In a horse race with the optimal model, balance matching has the best fit for slightly more than two-thirds of observations. When compared with the other heuristic models, balance matching is also the best fit for approximately two-thirds of observations. Balance matching performs comparably to, or only a little worse than, the machine learning models, with balance matching exhibiting the best fit for 40.7% to 49.9% of observations.

To the extent that we think of the competing models as actually representing different models of individual decision-making, we would naturally expect the best-fit model to be persistent within individuals over time. Table 6 shows the within-person transition matrix for the best-fit model. The sample is restricted to individual × months where we observe repayment behavior for at least two months in a row. For this exercise, we allow the set J to encompass all of the candidate models, and we fix the uniformly distributed repayment to be constant within an individual over time.

The table shows that balance matching exhibits a high degree of persistence – both in absolute value and relative to the other models of repayment behavior. Among individuals whose repayments are best fit by the uniform model in a given month, 22.5% make repayments that are closest to the uniform model in the next month. This persistence likely reflects the fact that balances and repayments are sticky over time – if the uniform model happens to be accurate in a given month, and balances and payments are sticky, then the uniform model will mechanically be accurate in the next month as well.

The balance-matching model exhibits three-fold greater persistence than the uniform model. Among individuals whose repayments are closest to balance matching in a particular month, 83.4% make payments that are closest to balance matching in the next month. The high degree of persistence suggests that balance matching is more than a good statistical model but is actually capturing a stable feature of individual decision-making. The only other model that exhibits strong persistence is the $1/N$ rule, which again likely reflects that fact that $1/N$

[36] Since the machine learning models were tuned to minimize RMSE, it is natural for these models to perform relatively better when evaluated using RMSE (and other distance metrics) than when evaluated using this type of horse race analysis.

[37] We exclude a comparison of the balance-matching and the uniform model, since it was shown in Panel A.

repayments (or the tendency to round payments) is a stable feature of individual behavior.

Taken together, our goodness-of-fit analysis supports the view that balance matching is a powerful predictor of credit card repayments, capturing more than half of the predictable variation in repayment behavior and performing substantially better than the alternative models. In the horse race analysis, balance matching performs at a similar level to the machine learning models, and is highly persistent over time, suggesting it is more than a good statistical model but is actually capturing a stable feature of individual decision-making.

7 Sensitivity Analysis

7.1 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 4). 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.[38]

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 (the “floor”) and a percentage of the balance (the “slope”). For instance, a typical minimum payment formula might be:

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

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

Consider the following scenarios for an individual with two cards:

(i) If minimum payments are on the "slope" part of the formula (balances greater than £1,250), and the slopes are identical (2% for both cards), then the balance-matching and minimum-payment-matching payments will be almost perfectly correlated.[39]

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

(iii) If minimum payments are on the "floor" 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 slopes (scenario ii) and where the floor binds (scenario iii) allows us to separately identify these mechanisms.

Figure 8 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 slope (scenario i), the middle row shows this relationship when the slopes are different (scenario ii), and the bottom row shows this relationship when both cards are on the floor part of the formula (scenario iii). The correlations between these different measures are shown in Table A7.

In the same slope 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 slope 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 floor 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 ($\rho = 0.23$).

[39] 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.

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 (2017), who estimate that 9% to 20% of account-holders anchor their repayments to minimum payment amounts.

7.2 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% 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% pay the minimum and 42.2% 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% 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 A6 plots repayment behavior for observations where individuals use autopay on both cards (left column, 11% of observations) and do not use autopay on either card (right column, 77% of observations).[40] The top row shows the distributions of actual and optimal payments, the middle row shows the distribution of actual and optimal payments in excess of

[40] The propensity to use autopay is highly correlated within individuals across cards. In the two-card sample, 68.2% of individuals use autopay on the high APR card also use it on the low APR card, and 74.9% of individuals who do not use autopay on the high APR card do not use it on the low APR card.

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% versus 23.2%), misallocated repayments in excess of the minimum are similar in both samples (45.7% versus 45.5%). 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.[41] Summary statistics for actual and excess payments by autopay status are shown in Appendix Table A8.

Appendix Table A9 and Table A10 show our standard measures of model performance by whether individuals use autopay on both cards and do not use autopay on either card.[42]. In particular, Appendix Table A9 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 A10 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.

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

[42] Results are shown using the 20% hold-out sample

8 Conclusion

In this paper, we used linked data on multiple cards from five major credit card issuers in the U.K. to study borrowing behavior in the credit card market. We showed that the allocation of repayments is highly non-optimal, with individuals allocating only 51.5% of their payments in excess of the minimum to the high APR card, relative to optimal repayments of 97.1%. This finding builds on Ponce et al. (2017), who showed evidence of similar non-optimal behavior in credit card data from Mexico.

The main contribution of our paper was to show that, in contrast to the optimal repayment rule, actual repayment behavior can be explained by a *balance-matching* heuristic under which individuals match the share of repayments on each card to the share of balances on each card. In particular, we showed that balance matching captures more than half of the predictable variation in repayments, performs substantially better than other models, and is highly persistent within individuals over time.

We provided additional support for the importance of balances – and irrelevance of interest rates – using machine learning models. Consistent with the poor fit of the optimal repayment rule, we found that interest rates have the lowest variable importance in our machine learning models. Consistent with balance matching results, we found that balances have the highest variable importance, with importance factors substantially larger than any of the other explanatory variables.

Our findings suggest a number of avenues for future work. Is balance-matching type behavior present in other markets, such as the decision of how to allocate payments across student loans with different interest rates? What are the psychological underpinnings of balance matching? For instance, does balance matching, in part, reflect a moral preference to repay more to the creditor to whom more is owed?[43] From a policy perspective, a natural response to the observed non-optimality of repayments is to make interest rates more salient. In our view, running experiments that test alternative ways of presenting interest rates on credit card statements and online repayment portals would be valuable.

[43] Devin Pope deserves credit for suggesting this underlying explanation.

References

Agarwal, S., S. Chomsisengphet, N. Mahoney, and J. Stroebel (2015). Regulating Consumer Financial Products: Evidence from Credit Cards. *Quarterly Journal of Economics 130*(1), 111–164.

Andersen, S., J. Campbell, K. M. Nielsen, and T. Ramadorai (2017). Inattention and Inertia in Household Finance: Evidence from the Danish Mortgage Market. *NBER Working Paper 21386*.

Ariely, D., G. Loewenstein, and D. Prelec (2003). "Coherent Arbitrariness": Stable Demand Curves Without Stable Preferences. *The Quarterly Journal of Economics 118*(1), 73–106.

Benartzi, S. and R. H. Thaler (2001). Naive Diversification Strategies in Defined Contribution Saving Plans. *American Economic Review 91*(1), 79–98.

Bergman, O., T. Ellingsen, M. Johannesson, and C. Svensson (2010). Anchoring and Cognitive Ability. *Economics Letters 107*(1), 66–68.

Breiman, L. (2001). Random Forests. *Machine learning 45*(1), 5–32.

Calvo, G. A. (1983). Staggered Prices in A Utility-Maximizing Framework. *Journal of Monetary Economics 12*(3), 383–398.

Campbell, J. Y. and J. F. Cocco (2003). Household Risk Management and Optimal Mortgage Choice. *The Quarterly Journal of Economics 118*(4), 1449–1494.

Chetty, R., J. N. Friedman, and J. E. Rockoff (2014). Measuring the Impacts of Teachers I: Evaluating Bias in Teacher Value-Added Estimates. *American Economic Review 104*(9), 2593–2632.

DellaVigna, S. (2009). Psychology and Economics: Evidence from the Field. *Journal of Economic Literature 47*(2), 315–372.

FCA (2016). Credit Card Market Study Final Report: Annex 2 Further Analysis. Technical report, Financial Conduct Authority.

Friedman, J., T. Hastie, and R. Tibshirani (2001). *The Elements of Statistical Learning*, Volume 1. Springer Series in Statistics New York.

Friedman, J. H. (2001, October). Greedy Function Approximation: A Gradient Boosting Machine. *Annals of statistics 29*(5), 1189–1232.

Herrnstein, R. J. (1961). Relative and Absolute Strength of Response as A Function of Frequency of Reinforcement. *Journal of the Experimental Analysis of Behavior 4*(3), 267–272.

Hortaçsu, A. and C. Syverson (2004). Product Differentiation, Search Costs, and Competition in the Mutual Fund Industry: A Case Study of S&P 500 Index Funds. *The Quarterly Journal of Economics 119*(2), 403–456.

Keys, B. J., D. G. Pope, and J. C. Pope (2016). Failure to Refinance. *Journal of Financial Economics 122*(3), 482–499.

Keys, B. J. and J. Wang (2017). Minimum Payments and Debt Paydown in Consumer Credit Cards. *NBER Working Paper*.

Mullianathan, S. and J. Spiess (2017). Machine Learning: An Applied Econometric Approach. *Journal of Economic Perspectives 31*, 87–106.

Ponce, A., E. Seira, and G. Zamarripa (2017). Borrowing on the Wrong Credit Card: Evidence from Mexico. *The American Economic Review 107*(4), 1335–1361.

Rubinstein, A. (2002). Irrational Diversification in Multiple Decision Problems. *European Economic Review 46*(8), 1369–1378.

Sims, C. A. (2003). Implications of Rational Inattention. *Journal of Monetary Economics 50*(3), 665–690.

Stango, V. and J. Zinman (2016). Borrowing High versus Borrowing Higher: Price Dispersion and Shopping Behavior in the U.S. Credit Card Market. *The Review of Financial Studies 29*(4), 979–1006.

Stewart, N. (2009). The Cost of Anchoring on Credit Card Minimum Payments. *Psychological Science 20*, 39–41.

Tversky, A. and D. Kahneman (1974). Judgment under Uncertainty: Heuristics and Biases. *Science 185*(4157), 1124–1131.

Vulkan, N. (2000). An Economist's Perspective on Probability Matching. *Journal of Economic Surveys 14*(1), 101–118.

Figure 1: Actual and Optimal Payments

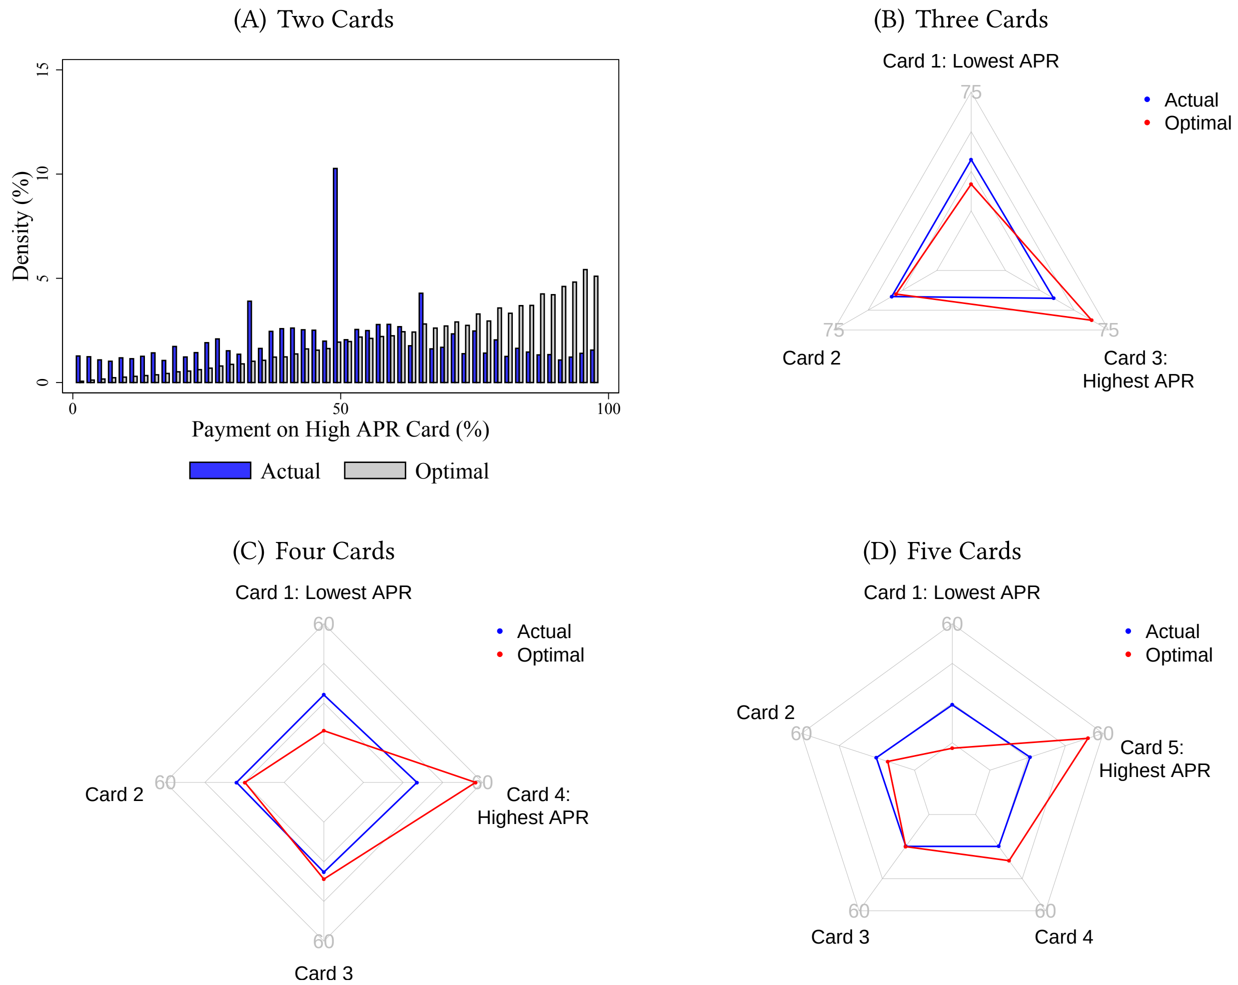

Note: Panel A shows the distribution of actual and optimal payments on the high interest rate card in the two-card sample. Panels B to D show radar plots of mean actual and optimal payments in the samples with 3 to 5 cards. In the radar plots, cards are ordered clockwise from the highest to the lowest 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 details on the sample construction.

Figure 2: Misallocated Payments by Economics Stakes

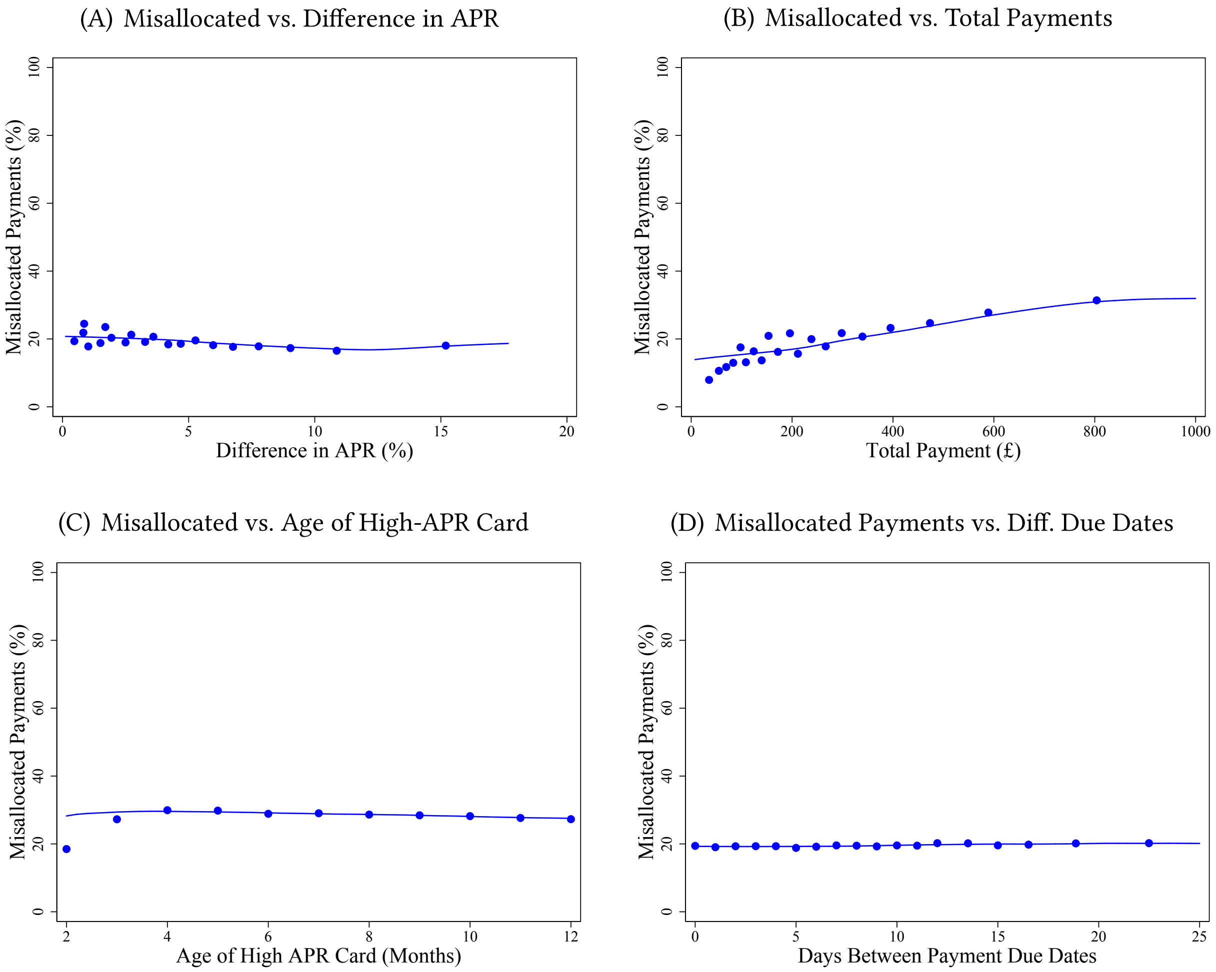

Note: Figure shows binned-scatter plots (with 20 equally sized bins) of misallocated payments against the difference in APR across cards (Panel A), the total value of payments within the month in pounds (Panel B), the age of the high APR card (Panel C), and the difference in payment due dates (Panel D). 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 3: Rounding and the $1/N$ Rule

(A) Density of Payments (£s)

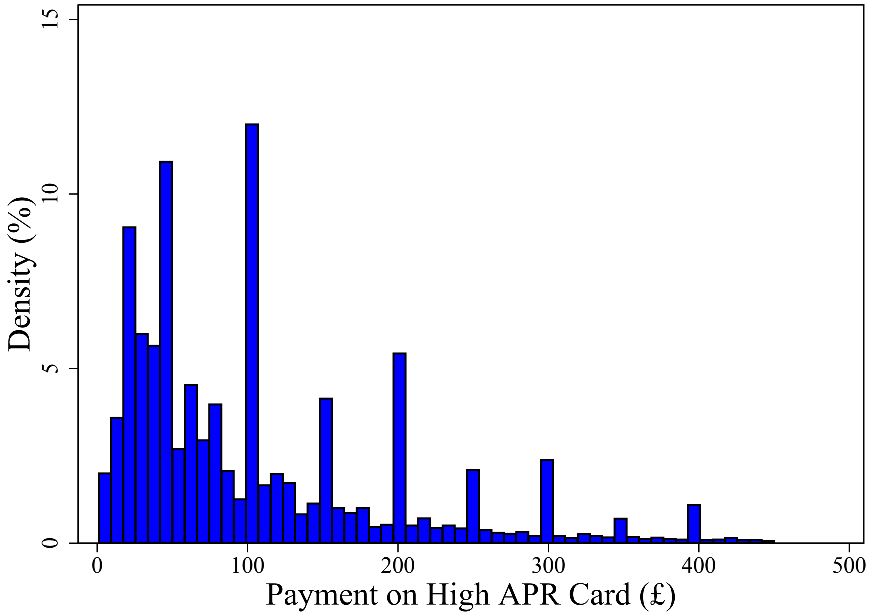

(B) Density of Payments (%), Round Number Values

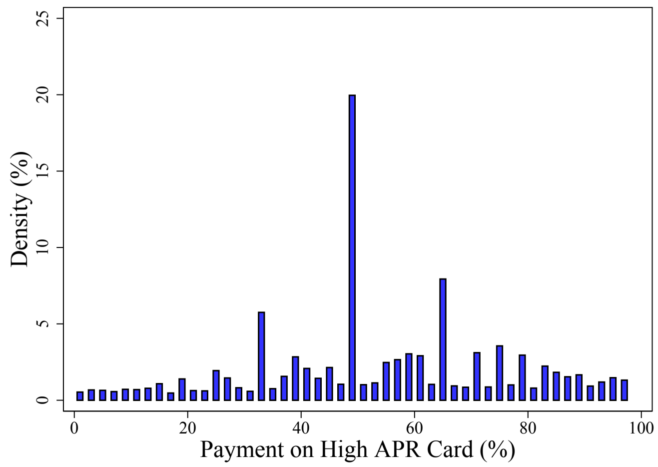

(C) Density of Payments (%), Non-Round Number Values

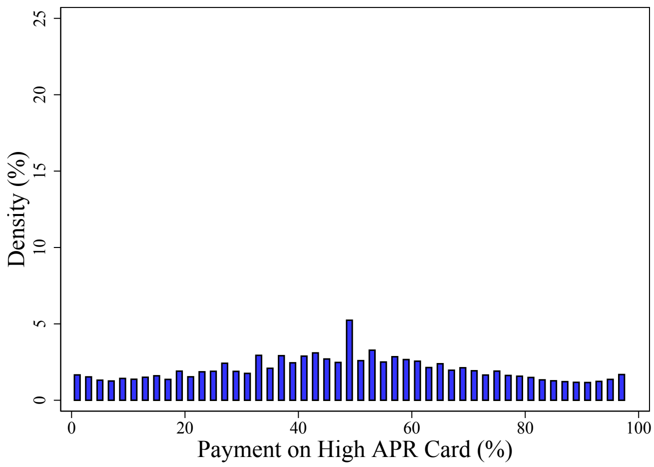

Note: Panel A shows the distribution of payments on the high APR card in £s (excluding the top decile of payments). Panel B plots the distribution of payments on the high APR card in percent, among individuals who make round number payments (exact multiples of £50). Panel C plots the distribution of payments on the high APR in percent, among individuals who make non-round number payments (not multiples of £50). The round and non-round samples are defined by repayments on the high APR card. See Footnote 27 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 4: Example Credit Card Statement

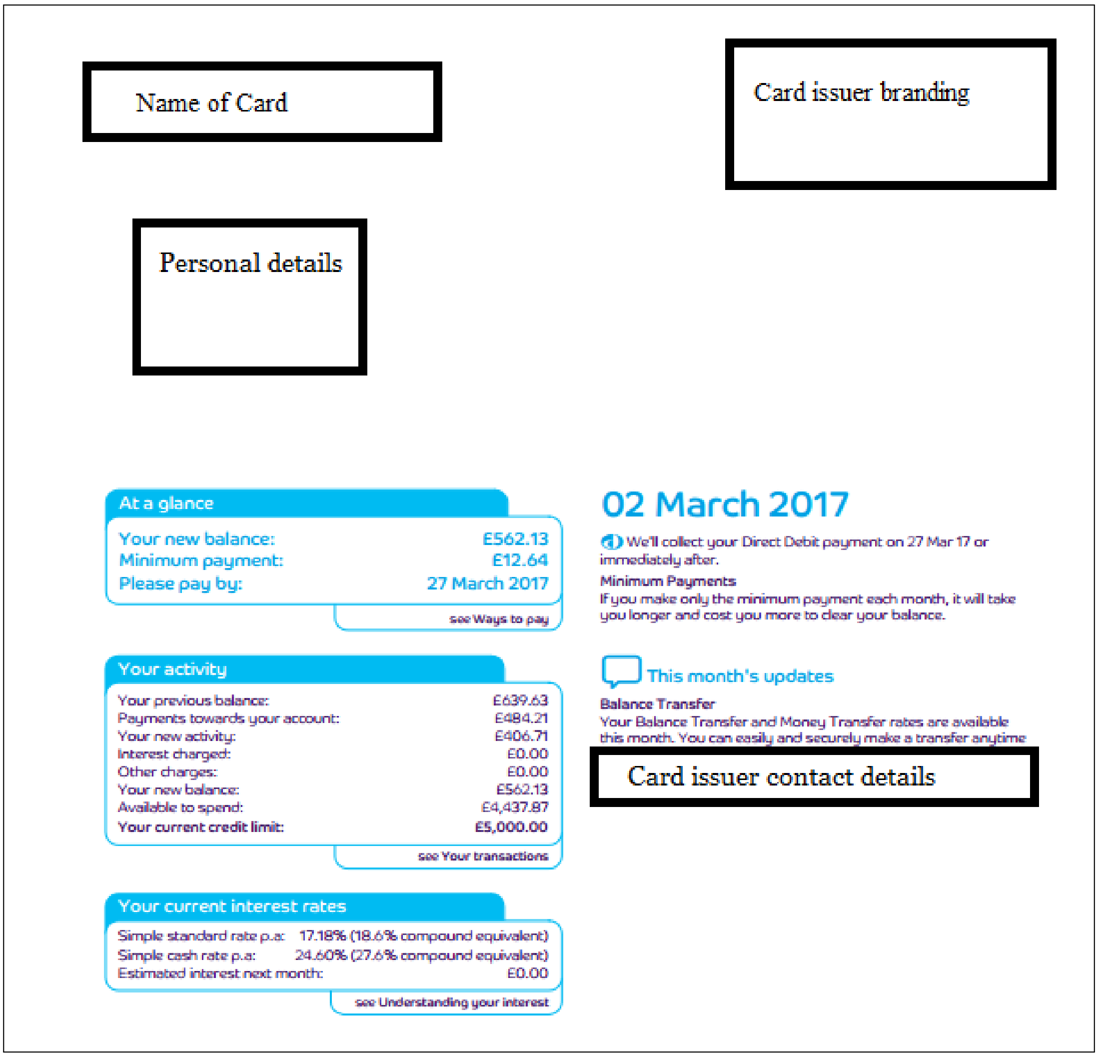

Name of Card

Card issuer branding

Personal details

At a glance

Your new balance:	£562.13
Minimum payment:	£12.64
Please pay by:	27 March 2017

see Ways to pay

Your activity

Your previous balance:	£639.63
Payments towards your account:	£484.21
Your new activity:	£406.71
Interest charged:	£0.00
Other charges:	£0.00
Your new balance:	£562.13
Available to spend:	£4,437.87
Your current credit limit:	**£5,000.00**

see Your transactions

Your current interest rates

Simple standard rate p.a:	17.18% (18.6% compound equivalent)
Simple cash rate p.a:	24.60% (27.6% compound equivalent)
Estimated interest next month:	£0.00

see Understanding your interest

02 March 2017

We'll collect your Direct Debit payment on 27 Mar 17 or immediately after.

Minimum Payments

If you make only the minimum payment each month, it will take you longer and cost you more to clear your balance.

This month's updates

Balance Transfer

Your Balance Transfer and Money Transfer rates are available this month. You can easily and securely make a transfer anytime

Card issuer contact details

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 5: Balance Matching

(A) Baseline Sample

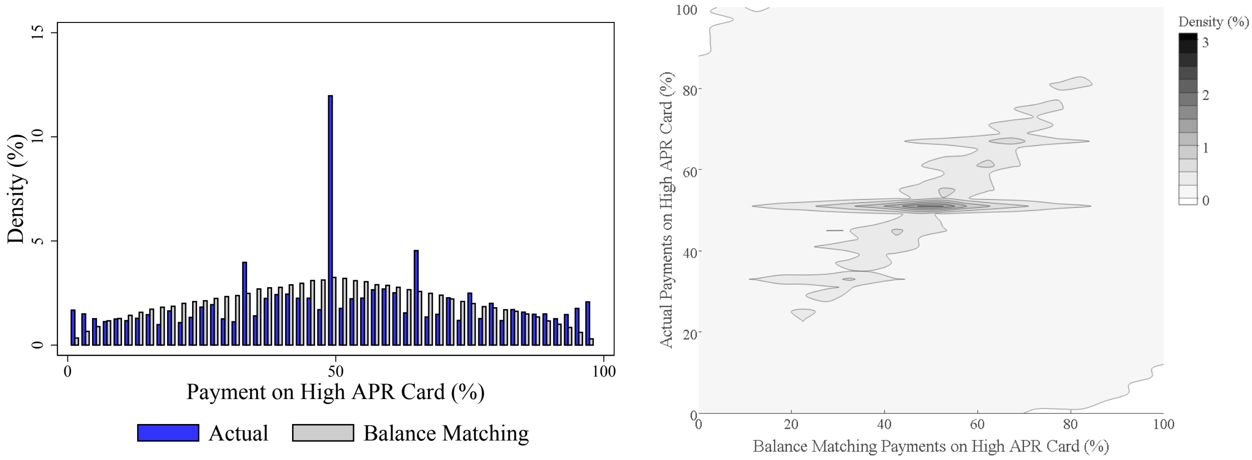

(B) Round Number Payment Sample

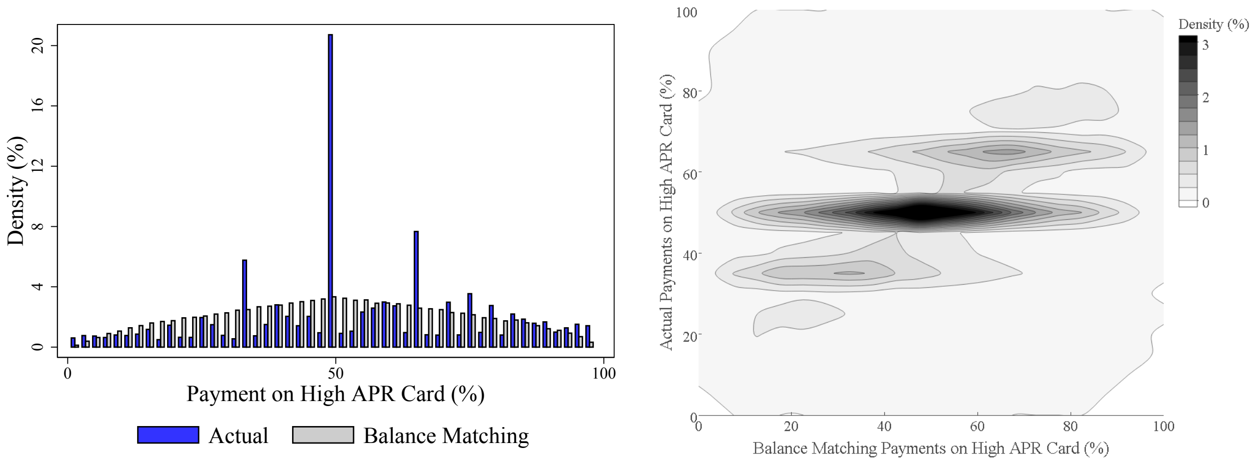

(C) Non-Round Number Payment Sample

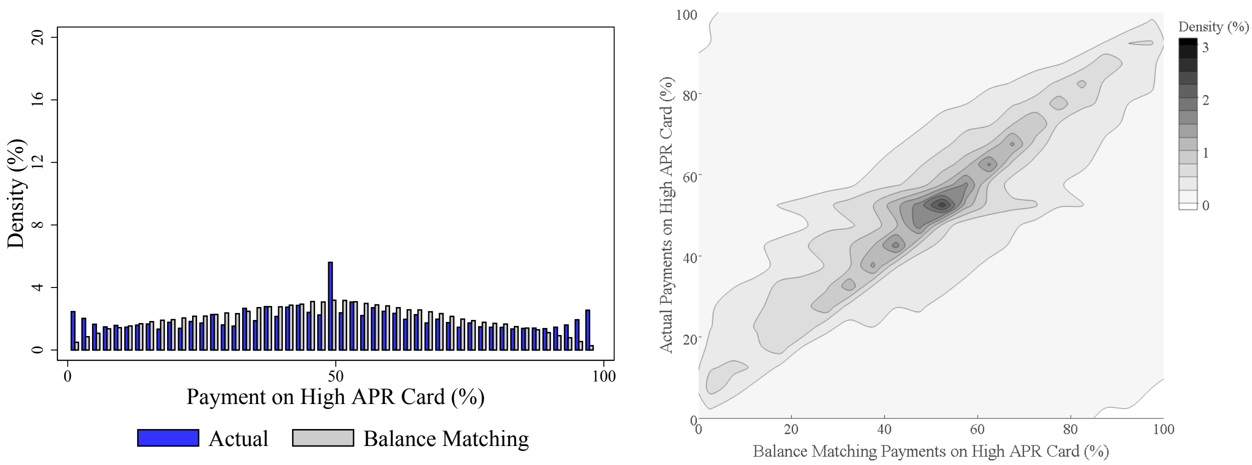

Note: Left panels shows the distribution of actual and balance-matching payments on the high APR card. Right panels show the joint density of actual and balance-matching payments. Panel A shows the baseline sample, Panel B restricts the sample to round number payments (multiples of £50), and Panel C restricts the sample to non-round payment amounts (not multiples of £50). Round and non-round samples are defined by repayments on the high APR card. See Footnote 27 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 6: Distribution of Actual and Balance-Matching Payments on Multiple Cards

(A) Three Cards

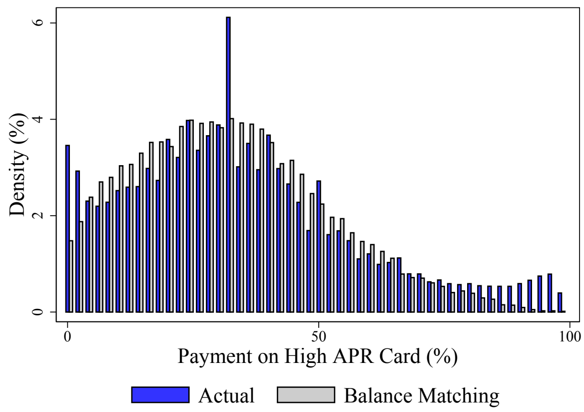

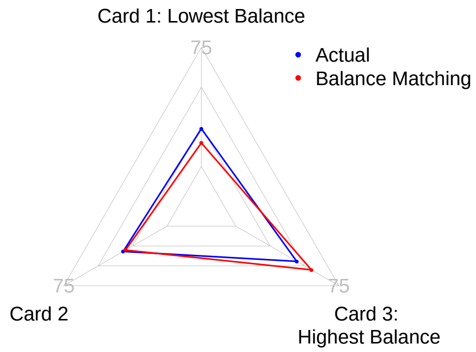

(B) Four Cards

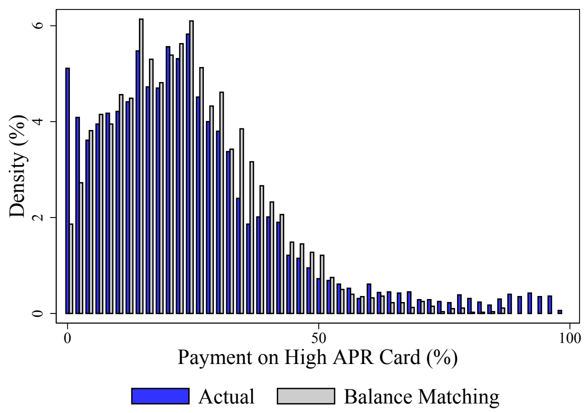

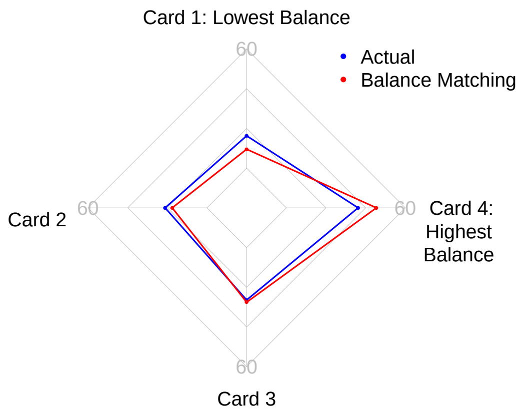

(C) Five Cards

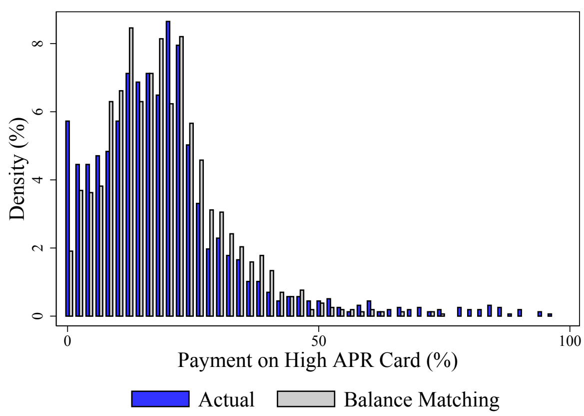

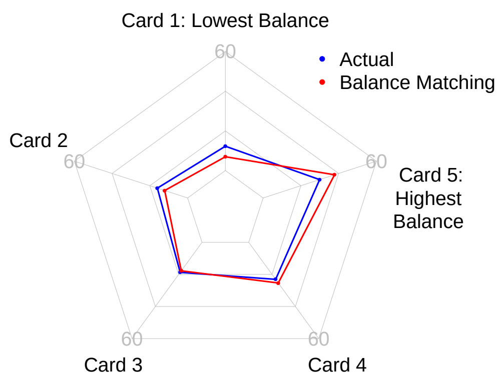

Note: Left column shows the marginal distributions of actual and balance-matching payments on the high APR card. Right column shows radar plots of the mean percentage of actual and balance-matching payments allocated to each card. In the radar plots, cards are ordered clockwise from the highest to the lowest balance (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 details on the sample construction.

Figure 7: Goodness-of-Fit for Different Models

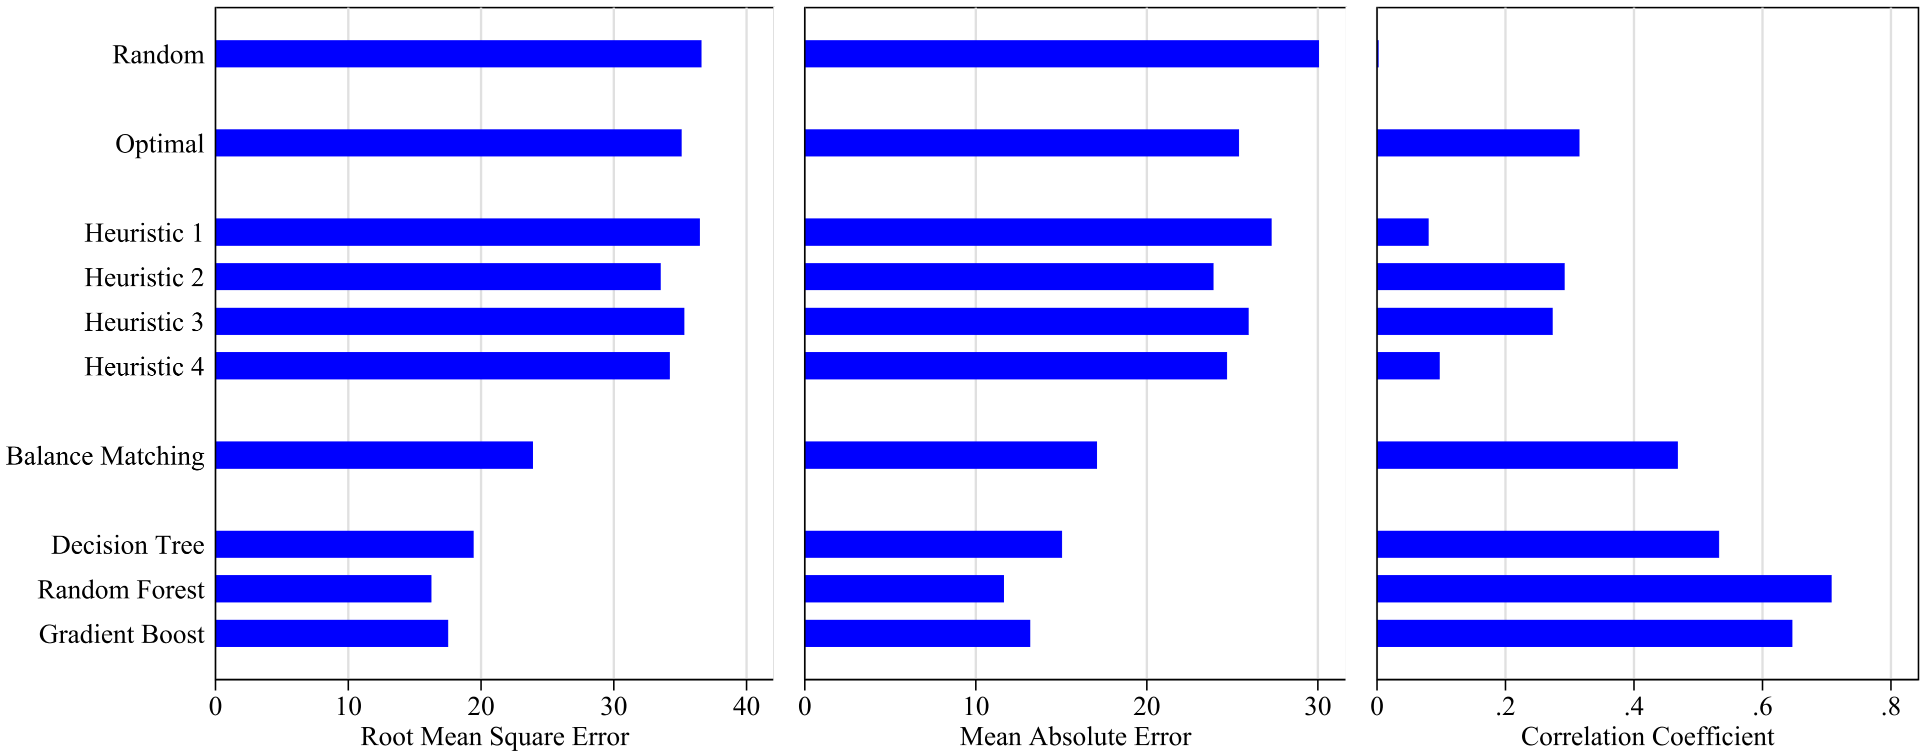

Note: Goodness-of-fit for different models of the percentage of payments on the high APR card. The left panel shows the Root Mean Square Error (RMSE), the middle panel shows the Mean Absolute Error (MAE), and the right panel shows the Pearson Correlation Coefficient, which can also be interpreted as the square root of the R-squared. Random has repayments on the high APR card randomly drawn from a uniform distribution with support on the 0% to 100% 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 the 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. 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% hold-out sample. See Section 2.2 for details on the sample construction.

Figure 8: Balance Matching and Minimum Payment Matching in the Floor and Slope Samples

(A) Same Slope Sample

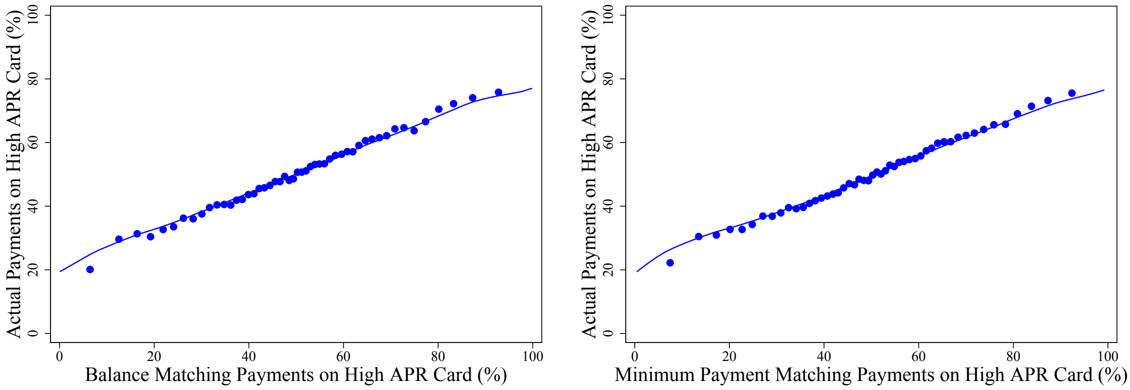

(B) Different Slope Sample

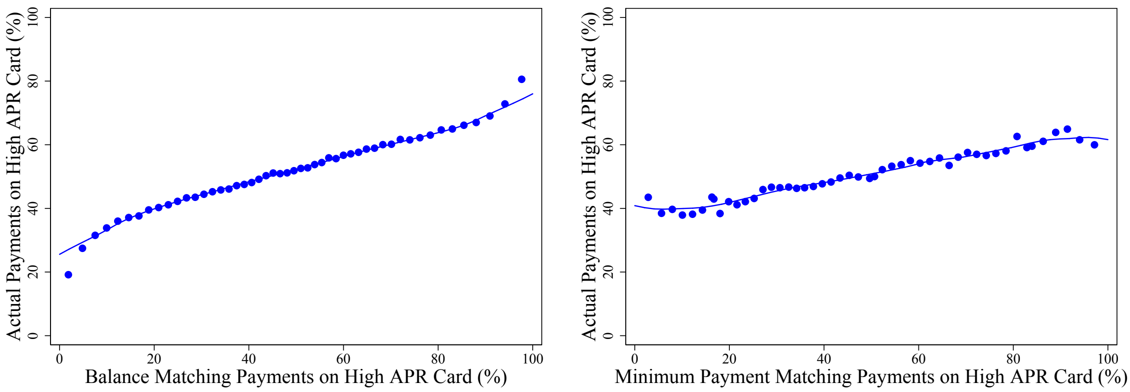

(C) Floor Sample

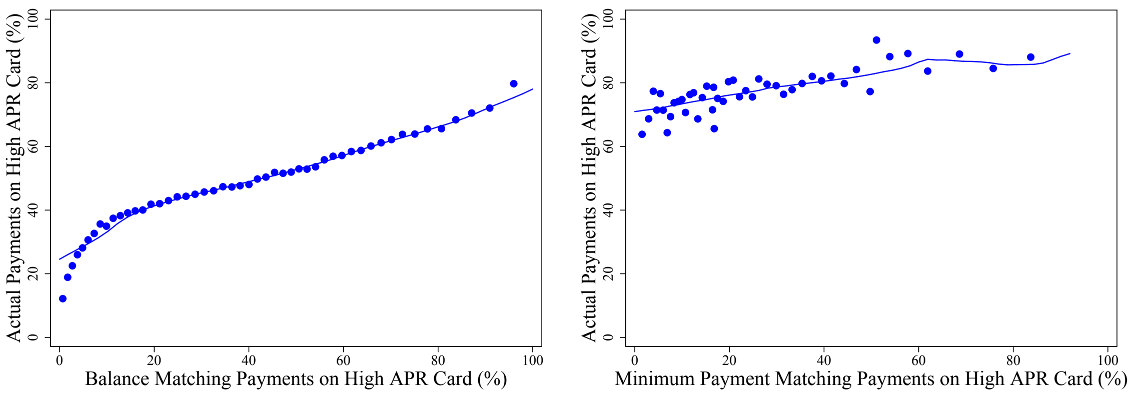

Note: Panels show binned-scatter plots of the actual percentage of monthly payment allocated to the high APR card (y-axis) and the percentage of total monthly payment allocated to the high APR card under the balance-matching heuristic (x-axis, left column) and minimum-payment-matching heuristics (x-axis,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

Table 1: Summary Statistics

	(1) High APR Card		(2) Low APR Card		(3) Difference	
	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Card Characteristics						
APR: Purchases (%)	22.87	4.80	16.56	6.40	6.30	5.85
APR: Cash Advances (%)	26.08	4.12	23.72	5.28	2.36	6.31
Monthly Credit Limit (£)	6,388.77	4,443.05	6,013.20	4,092.41	375.57	4,856.48
Spending (£)						
Purchases	128.09	432.43	116.53	397.63	11.56	570.04
Purchases if > £0	380.17	672.79	360.06	629.70	-2.80	798.02
Cash Advances	6.47	73.29	5.81	73.74	0.66	97.25
Cash Advances if > £0	216.98	366.68	215.01	395.20	-5.42	352.15
Payments (£)						
Repayments	259.76	733.92	229.69	657.60	30.07	913.65
Interest Paid (£)						
Purchases	38.48	59.49	28.97	48.32	9.51	61.64
Cash Advances	1.49	10.73	0.91	7.13	0.58	11.88
Card Cycle (£)						
Closing Balance	3,020.54	3,115.48	3,032.15	2,967.13	-11.61	3,478.14
Revolving Balance	2,200.01	2,890.49	2,053.68	2,796.17	146.33	3,082.07
Minimum Amount Due	63.24	68.84	56.80	58.32	6.43	71.55
Card Status						
Predicted Account Charge-Off Rate (%)	1.80	3.03	1.65	2.56	0.13	3.11
Tenure (Months Since Account Opened)	104.82	78.13	78.53	70.10	26.30	84.55
Number of Account-Months	394,111		394,111		394,111	

Note: Summary statistics for the two-card sample restricted to individual × months in which individuals face an economically meaningful allocative decision. See Section 2.2 for more details on the sample construction. APR stands for annual percentage rate. Predicted charge-off rate is the predicted probability that the credit card is charged off within the next six months. The exchange rate was £1 = \$1.32 at the midpoint of our sample period.

Table 2: Actual and Optimal Payments on the High APR Card

	Mean	Std. Dev.	Percentiles				
			10th	25th	50th	75th	90th
i) As % Total Monthly Payment							
Actual Payment (%)	51.22	24.21	16.86	33.33	50.00	67.99	84.78
Optimal Payment (%)	70.74	22.17	38.10	55.92	75.23	89.48	95.83
Difference (%)	19.52	23.75	0.00	0.72	9.91	32.40	54.55
ii) Payment in £							
Actual Payment (£)	259.76	733.92	25.00	45.49	100.00	200.00	450.00
Optimal Payment (£)	377.30	849.70	32.62	65.00	138.51	307.09	807.21
Difference (£)	117.54	422.14	0.00	1.00	17.80	75.00	237.47

Note: Summary statistics for actual and optimal payments on the high APR card. The top panel shows values as a percentage of total 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 more details on the sample construction. The exchange rate was £1 = $1.32 at the midpoint of our sample period.

Table 3: Annualized Interest Savings from Optimizing Credit Card Repayments (£s)

	Mean	Std. Dev.	Percentiles	
			75th	90th
Two Cards	64.82	115.33	70.39	167.41
Three Cards	121.26	463.63	133.44	414.36
Four Cards	198.40	665.57	262.80	703.68
Five Cards	247.65	851.83	366.96	926.88

Note: 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. Samples are restricted to individual $\times$ 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 4: 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 5: Horse Races Between Alternative Models

Panel (A)
Uniform vs. Other Rules

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Win %									
Uniform	32.81	45.45	50.24	45.28	47.41	46.87	30.70	23.43	27.23
Balance Matching	67.19								
Optimal		54.55							
Heuristic 1 (Pay Down Lowest Capacity)			49.76						
Heuristic 2 (Pay Down Highest Capacity)				54.72					
Heuristic 3 (Pay Down Highest Balance)					52.59				
Heuristic 4 (Pay Down Lowest Balance)						53.13			
Decision Tree							69.30		
Random Forest								76.57	
XGB									72.77

Panel (B)
Balance Matching vs. Other Rules

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Win %								
Balance Matching	68.97	73.09	67.22	76.65	64.70	49.86	40.74	45.33
Optimal	31.03							
Heuristic 1 (Pay Down Lowest Capacity)		26.91						
Heuristic 2 (Pay Down Highest Capacity)			32.78					
Heuristic 3 (Pay Down Highest Balance)				23.35				
Heuristic 4 (Pay Down Lowest Balance)					35.30			
Decision Tree						50.14		
Random Forest							59.26	
XGB								54.67

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 each of our models one-by-one 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% to 100% interval. Panel B compares each of the models one-by-one to the balance-matching model. We exclude a comparison of balance matching and the uniform model, since it was shown in Panel A. Uniform has repayments on the high APR card randomly drawn from a uniform distribution with support on the 0% to 100% interval. Balance matching is match the share of repayments on each card to the share of balances on each card. 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"). Decision Tree, Random Forest, and XGB are machine learning models that predict the share of repayments on the high APR card using these methods. 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% hold-out sample. See Section 2.2 for more details on the sample construction.

Table 6: Transition Matrix for Best-Fit Model

Previous Period	*Current Period*							
	Uniform	Balance Matching	Optimal	H 1	H 2	H 3	H 4	$1/N$ Rule
Uniform	25.50%	55.20%	0.23%	1.75%	0.47%	1.52%	0.58%	14.74%
Balance Matching	7.73%	83.35%	0.13%	1.39%	1.20%	1.17%	1.10%	3.92%
Optimal	18.18%	72.73%	9.09%	0.00%	0.00%	0.00%	0.00%	0.00%
Heuristic 1 (Pay Down Lowest Capacity)	11.38%	45.51%	0.00%	29.94%	0.00%	0.60%	4.19%	8.38%
Heuristic 2 (Pay Down Highest Capacity)	4.20%	66.39%	0.00%	0.84%	24.37%	1.68%	0.00%	2.52%
Heuristic 3 (Pay Down Highest Balance)	11.43%	47.62%	0.00%	4.76%	0.00%	19.05%	0.00%	17.14%
Heuristic 4 (Pay Down Lowest Balance)	4.55%	63.64%	0.00%	4.55%	0.00%	1.52%	21.21%	4.55%
$1/N$ Rule	17.58%	27.73%	0.26%	2.08%	0.39%	1.43%	0.52%	50.00%

Note: Table shows transition matrix for the best-fit payment model. The target variable is the share of repayments on the high APR card. In each period, we allow the set J to encompass all of the candidate models, and we fix the uniformly distributed repayment to be constant within an individual over time. Balance matching is match the share of repayments on each card to the share of balances on each card. Optimal is to 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"). $1/N$ is exactly split repayments across cards. Samples are restricted to individual $\times$ months in which individuals face an economically meaningful allocative decision. All results shown in the table are based on the 20% hold-out sample. See Section 2.2 for more details on the sample construction.

Online Appendix

I 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.[44] 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% of the data to train the tree, and calculate the error on each 20% subset.[45] Appendix Figure A7 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

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

[45] See Friedman et al. (2001) Chapter 9, for further information on tree-based methods.

(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.[46] 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 et al., 2001). Instead of growing a number of 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.[47] 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.[48]

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

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

[48] 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 et al. (2001), Chapter 10 also introduces a boosting algorithm.

Figure A1: Actual and Optimal *Excess* Payments

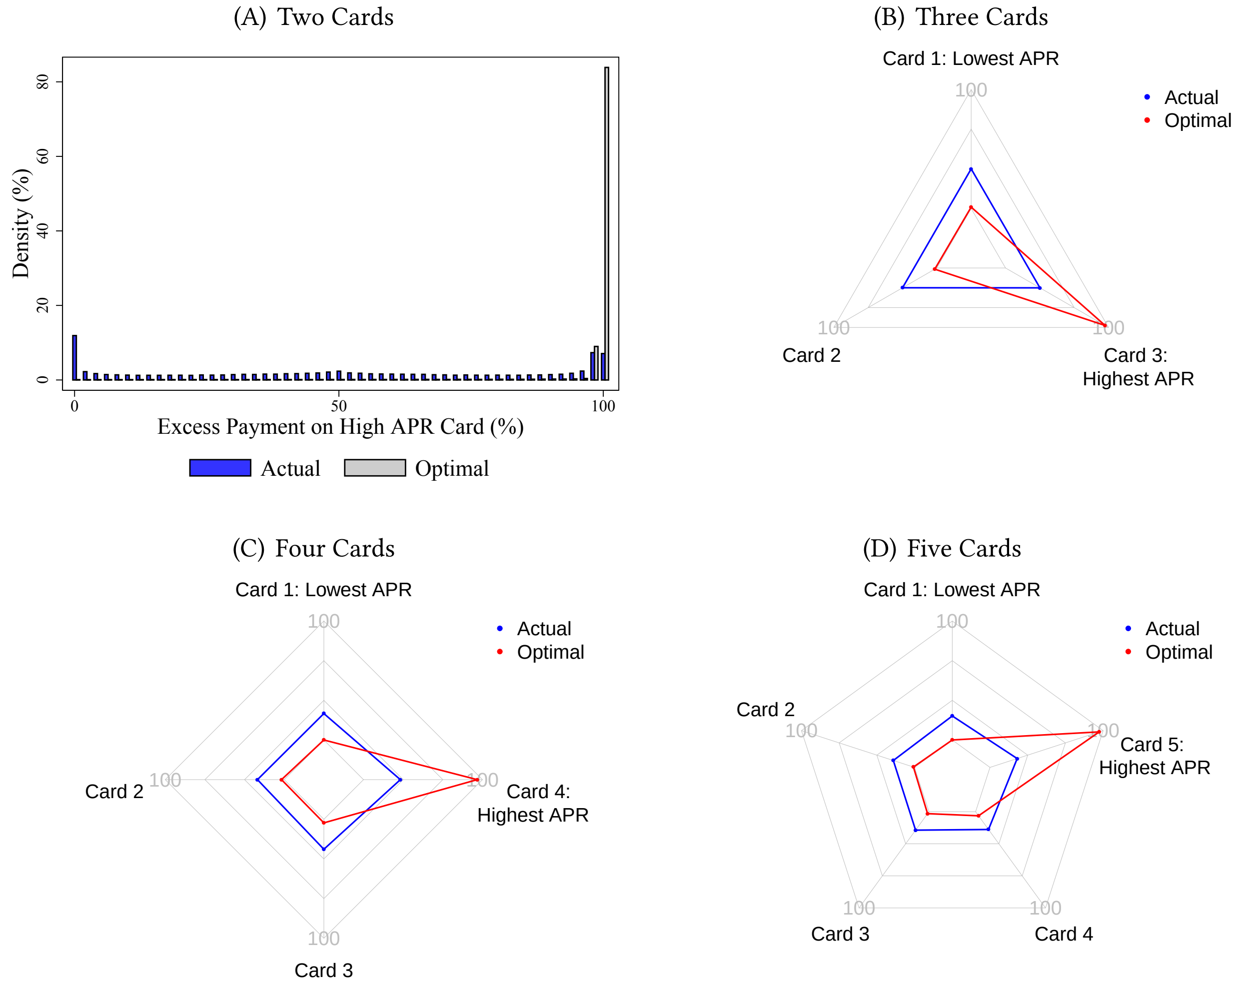

Note: Panel A shows the distribution of actual and optimal excess payments on the high interest rate card in the two-card sample. Panels B to D show radar plots of mean actual and optimal excess payments in the samples of individuals with 3 to 5 cards. Excess payments are calculated as the percentage 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

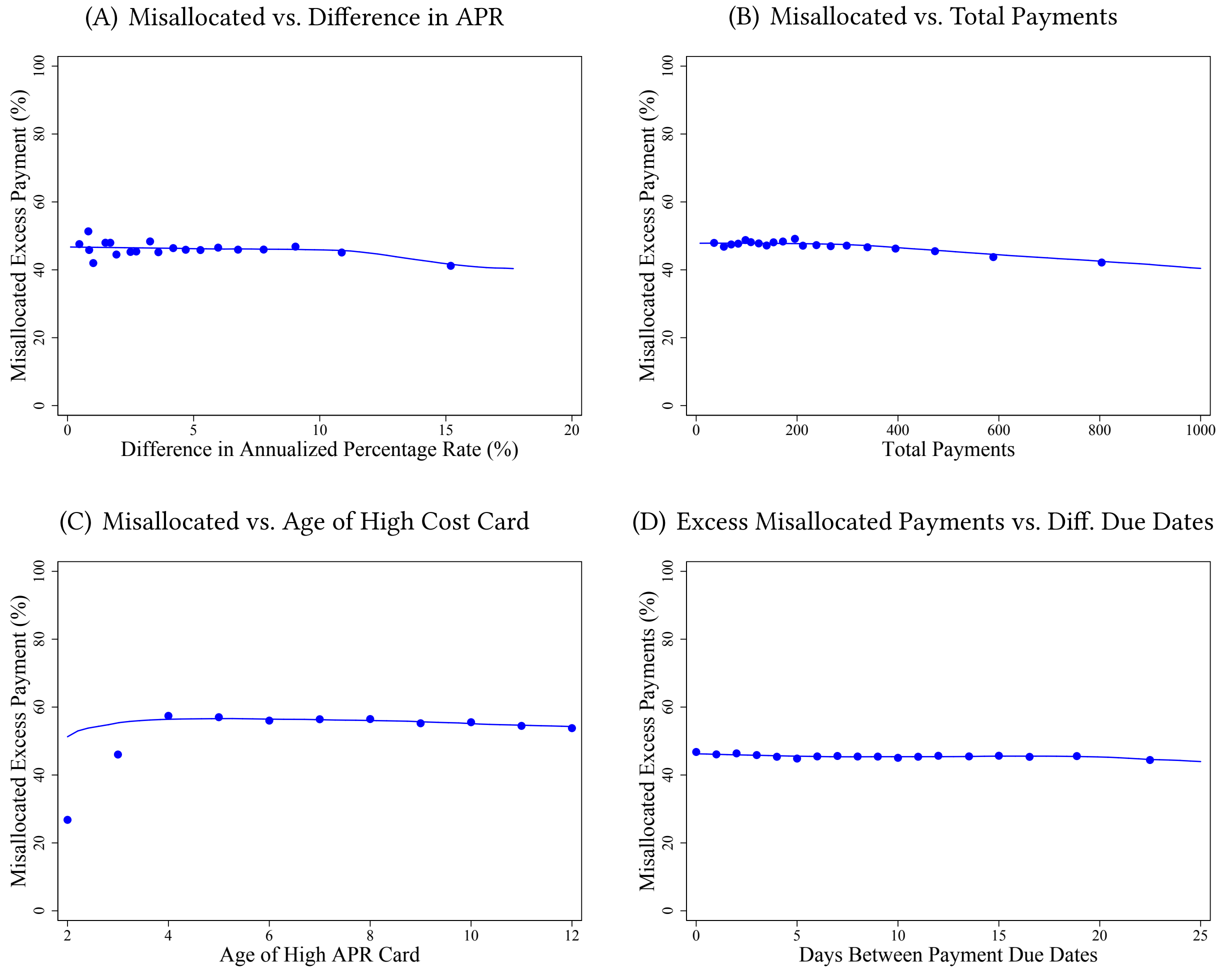

Note: Figure shows binned-scatter plots (with 20 equally sized bins) of misallocated payments against the difference in APR across cards (Panel A), the total value of payments within the month in pounds (Panel B), the age of the high APR card (Panel C), and the difference in payment due dates (Panel D). Local polynomial lines of best fit, based on the non-binned data, are also shown. 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 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 A3: Goodness-of-Fit for Different Models, Round and Non-Round Number Samples

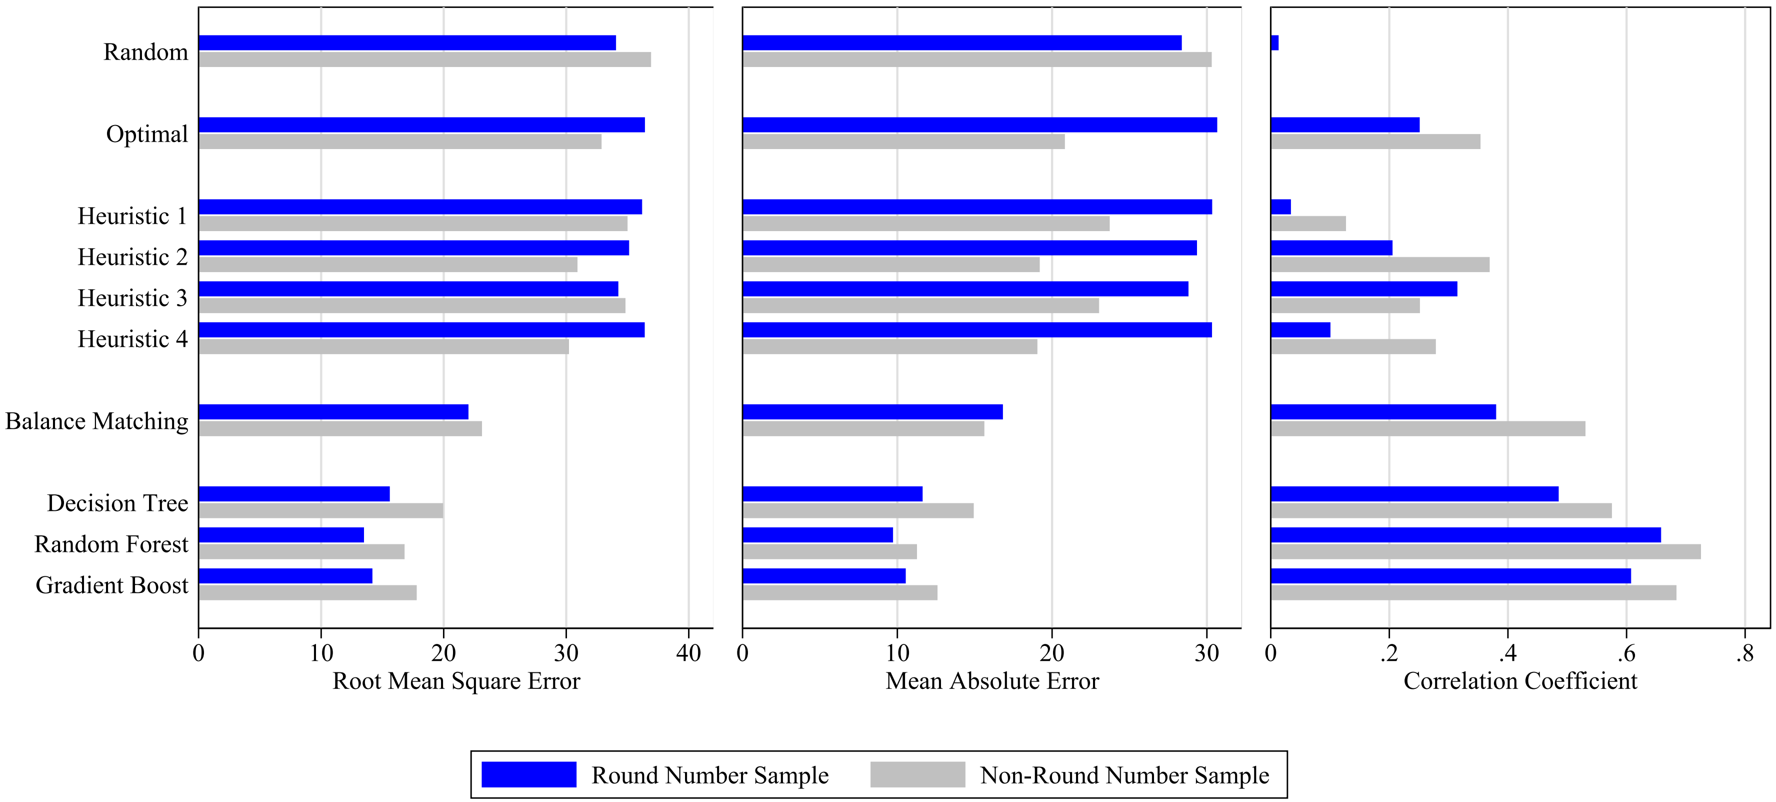

Note: Goodness-of-fit for different models of the percentage of payments on the high APR card. The left panel shows the Root Mean Square Error (RMSE), the middle panel shows the Mean Absolute Error (MAE), and right panel shows the Pearson Correlation Coefficient, which can also be interpreted as the 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% to 100% 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 27 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% hold-out sample. See Section 2.2 for more details on the sample construction.

Figure A4: Actual and Predicted Payments Under Alternative Repayment Heuristics

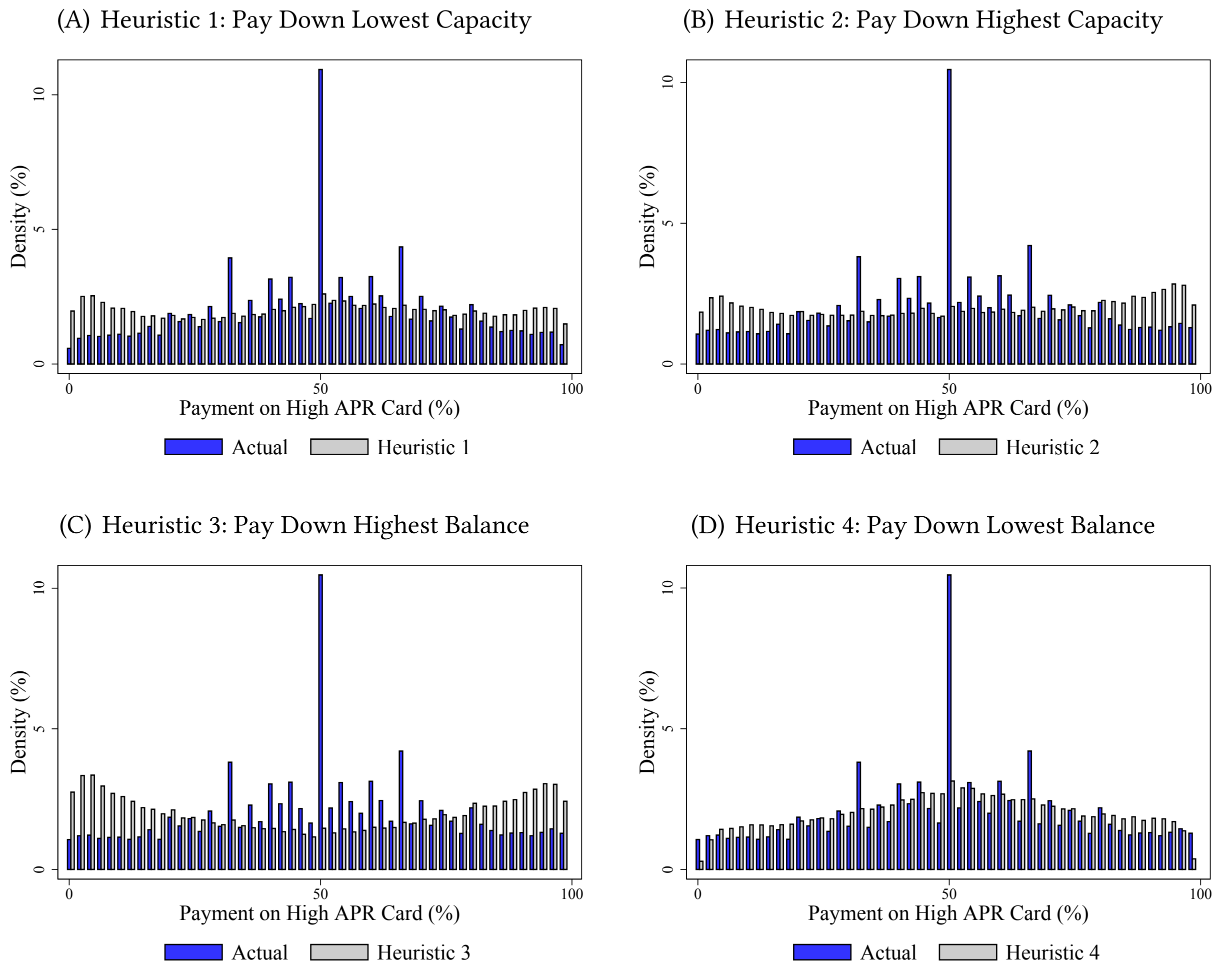

(A) Heuristic 1: Pay Down Lowest Capacity

(B) Heuristic 2: Pay Down Highest Capacity

(C) Heuristic 3: Pay Down Highest Balance

(D) Heuristic 4: Pay Down Lowest Balance

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 A5: Histogram of Difference in Due Dates

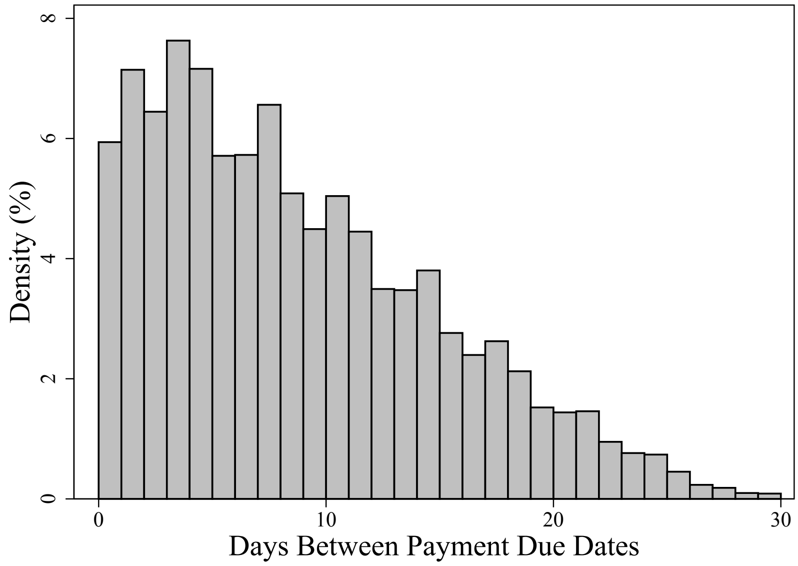

Note: Figure shows the distribution of the absolute difference in due dates. The two-card sample is restricted to individual × months in which individuals face an economically meaningful allocative decision. See Section 2.2 for more details on the sample construction.

Figure A6: Actual, Optimal and Balance Matching Payments for Autopay (Left Column, 11% of Observations) and Non-Autopay (Right Column, 77% of Observations) Samples

(A) Actual vs. Optimal Payments

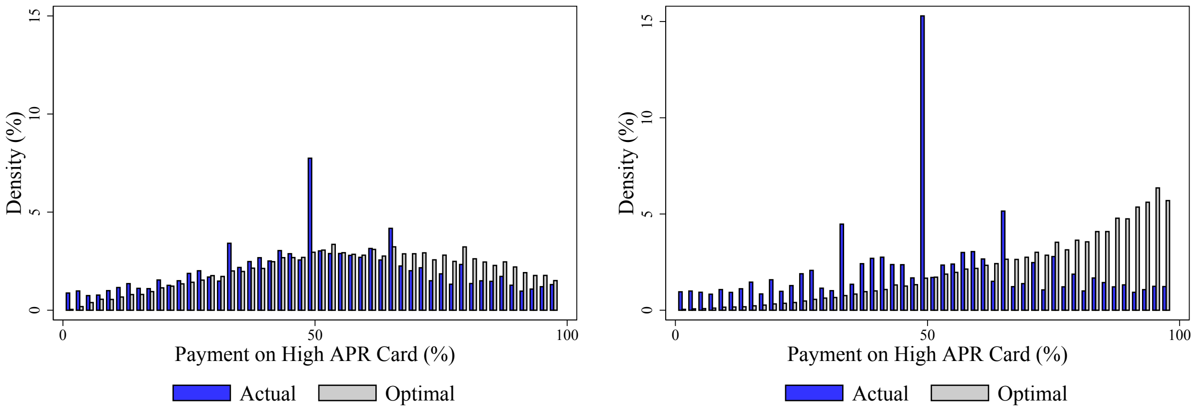

(B) Actual vs. Optimal Excess Payments

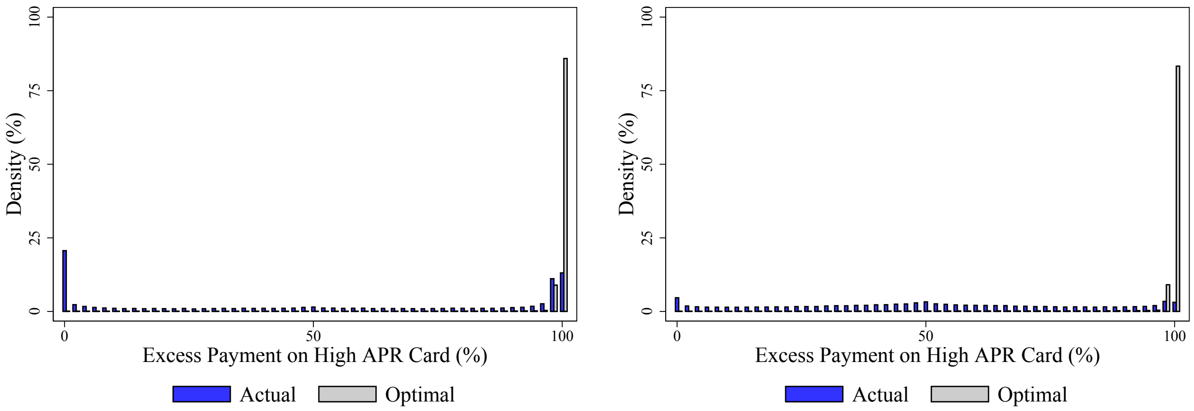

(C) Actual vs. Balance Matching Payments

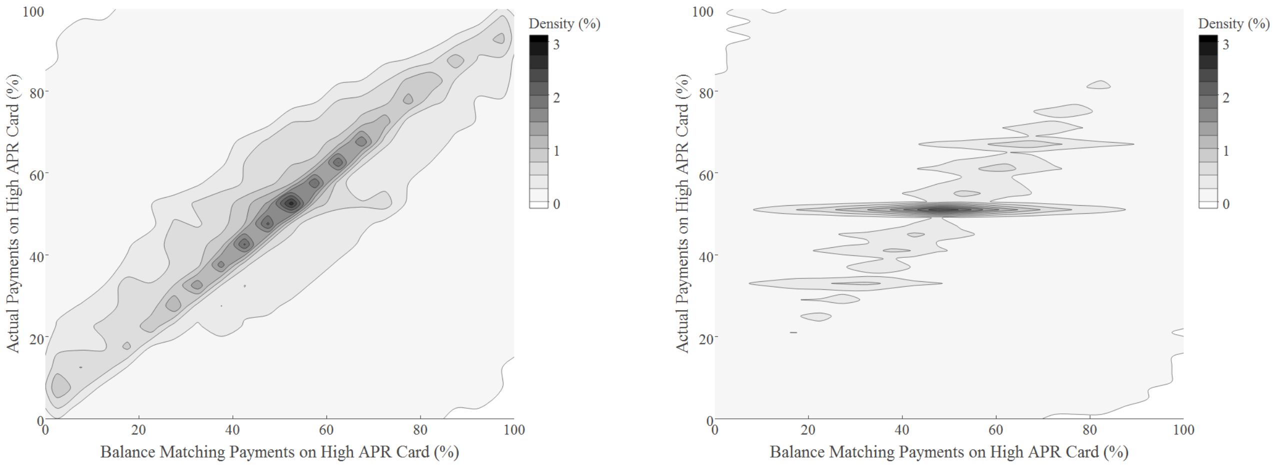

Note: Panel A shows the distribution of actual and optimal excess payments on the high APR card in the two-card sample. Panel B shows the distribution of actual and optimal excess payments on the high APR card in the two-card sample. Excess payments are calculated as the percentage of payments on a given 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.

Figure A7: High APR Card Payment Decision Tree

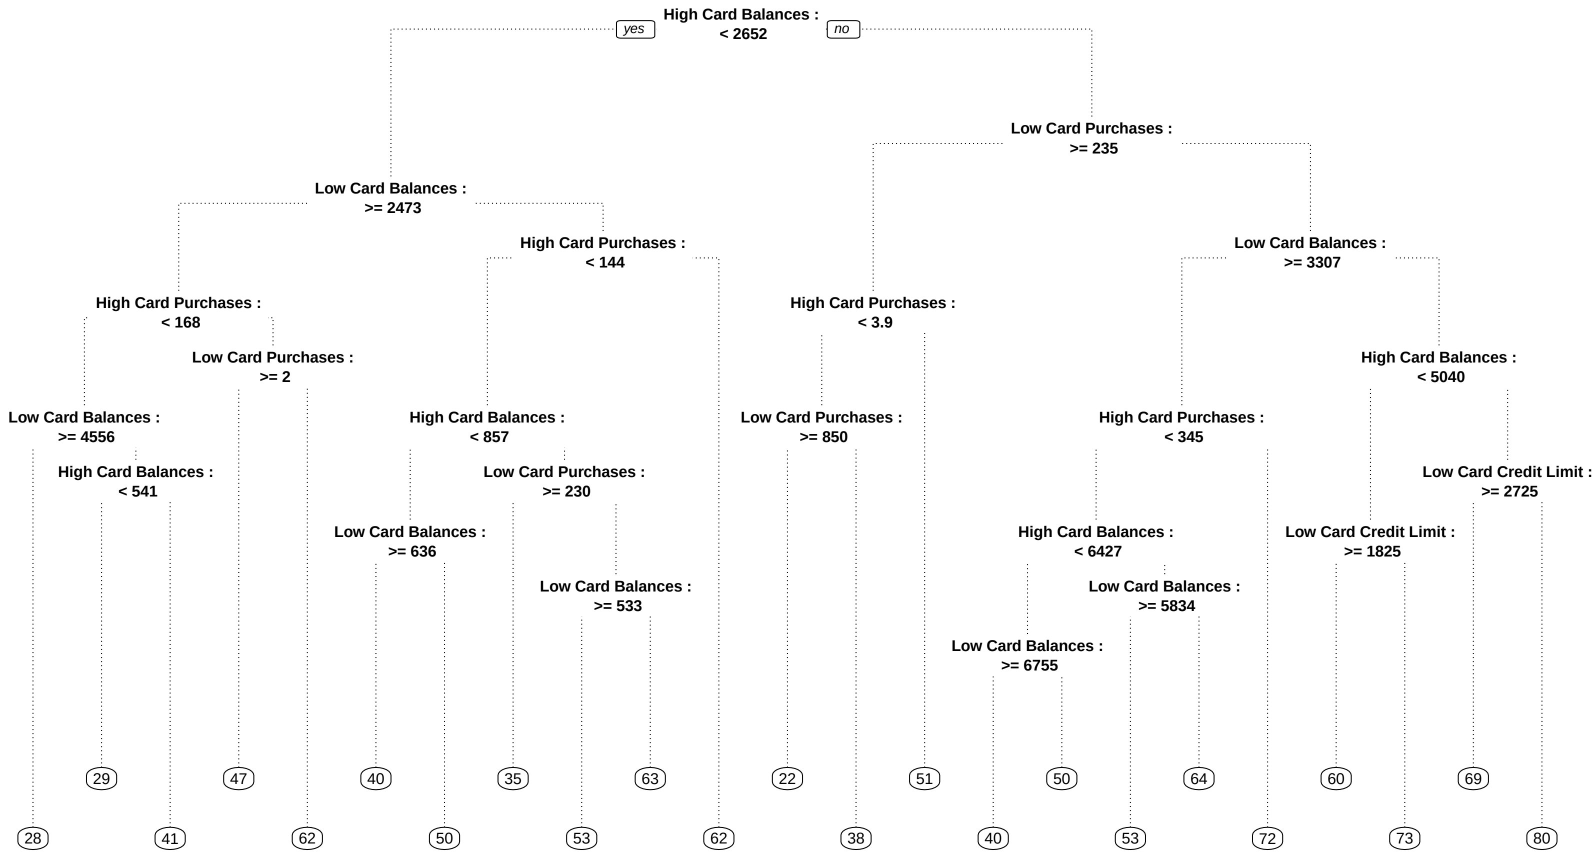

Note: Figure shows the decision (regression) tree for high APR card repayment. Top row is tree root. Nodes show the variable and split value at each branch. Bottom rows show predicted values at the end of each branch.

Table A1: Sample Restrictions

	(1) Unique Individuals		(2) Aggregate Revolving Debt	
	Count	%	£s	%
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

			Percentiles				
	Mean	Std. Dev.	10th	25th	50th	75th	90th
i) As a % Total Monthly Payment							
Actual Excess Payment (%)	51.51	34.75	0.89	19.92	51.31	84.91	99.82
Optimal Excess Payment (%)	97.08	12.93	100.00	100.00	100.00	100.00	100.00
Difference (%)	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: 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)
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 × months in which individuals face an economically meaningful allocative decision. Goodness of fit is calculated using the 20% hold-out sample and standard errors are constructed by the bootstrap method. See Section 2.2 for details.

Table A4: Goodness-of-Fit for Different Models, Round Number and Non-Round Number Payment Samples

	Non-Round Number Sample			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)
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 27 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% hold-out sample and standard errors are constructed by the bootstrap method. SSee Section 2.2 for details on the sample construction.

Table A5: 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 × months in which individuals face an economically meaningful allocative decision. See Section 2.2 for details on the sample construction.

Table A6: 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% training sample used in Table 4.

Table A7: Correlations Between Payment Rules

	Panel (A) Balance Matching vs. Min. Pay Matching		
	(1)	(2)	(3)
	Same Slopes	Different Slopes	Floor
Correlation	0.96	0.86	0.56
	(0.00)	(0.00)	(0.02)
	Panel (B) Balance Matching vs. Actual		
	(1)	(2)	(3)
	Same Slopes	Different Slopes	Floor
Correlation	0.63	0.41	0.50
	(0.00)	(0.00)	(0.00)
	Panel (C) Min. Pay Matching vs. Actual		
	(1)	(2)	(3)
	Same Slopes	Different Slopes	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 A8: Summary Statistics for Autopay (11% of Observations) and Non-Autopay (77% of Observations) Samples

	(1) Both Cards Non-Autopay	(2) Both Cards Autopay
i) Actual and Optimal Payments		
Actual Payments (%)	51.21	51.11
Optimal Payments (%)	74.36	58.42
Actual - Optimal Payments (%)	23.15	7.30
ii) Actual and Optimal Excess Payments		
Actual Excess Payments (%)	51.29	52.26
Optimal Excess Payments (%)	96.97	97.73
Actual Excess - Optimal Excess Payments (%)	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 $\times$ months in which individuals face an economically meaningful allocative decision. See Section 2.2 for details on the sample construction.

Table A9: 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)
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% of observations) is defined as observations where individuals make automatic payments on both cards. The non-autopay sample (77% of observations) 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. Goodness of fit is calculated using the 20% hold-out sample and standard errors are constructed by the bootstrap method. See Section 2.2 for details on the sample construction.

Table A10: 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 %		
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 %		
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% to 100% interval. Panel B compares optimal model repayments to the balance-matching model. The autopay sample (11% of observations) is defined as observations where individuals make automatic payments on both cards. The non-autopay sample (77% 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% hold-out sample. See Section 2.2 for more details on the sample construction.