

How Do Individuals Repay Their Debt?

The Balance-Matching Heuristic*

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

August 14, 2018

Abstract

We study how individuals repay their debt using linked data on multiple credit cards. Repayments are not allocated to the higher interest rate card, which would minimize the cost of borrowing. Moreover, the degree of misallocation is invariant to the economic stakes, which is inconsistent with optimization frictions. Instead, we show that repayments are consistent with a balance-matching heuristic under which the share of repayments on each card is matched to the share of balances on each card. Balance matching captures more than half of the predictable variation in repayments and is highly persistent within individuals over time.

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, Oslo 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, Leverhulme grant RP2012-V-022, and by the Becker Friedman Institute at the University of Chicago.

[†] 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, Warwick Business School; Network for Integrated Behavioural Science. Email: neil.stewart@wbs.ac.uk.

[¶] University of Nottingham, School of Economics; Network for Integrated Behavioural Science; Centre for 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 an important input for many fields of economic research and is directly relevant for consumer financial policy.

In this paper, we examine competing models of how individuals make debt payments – 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, holding the total repayment in each month fixed.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 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 to jointly estimate preference parameters.

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.

The first part of the paper documents that individuals do not optimally allocate their payments across cards. 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.⁵

This finding of non-optimal repayments closely matches prior research by Ponce et al. (2017) on credit card borrowing in Mexico. In particular, the most comparable result is that the share of payments in excess of the minimum misallocated to the high APR card is 50% among Mexican credit card holders and 46% among U.K. credit card holders. A possible explanation for the non-optimal repayments was that Mexican borrowers may not have had enough experience with credit cards, or enough financial sophistication, compared to borrowers in highly-developed countries. The results in our paper show, to the contrary, the same degree of non-optimal repayments even in the U.K., providing a striking example of uniformity of a behavioral bias across very different cultures and financial settings.⁶

The second part of the paper examines whether optimization frictions can explain the observed repayment behavior. In optimal inattention models, individuals face a fixed cost of optimization – such as the time, psychological, or cognitive costs associated with determining the optimal repayment allocation (Sims, 2003). Because of this fixed cost, these models predict

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 U.K. has a long history and high rate of credit card use. Credit cards were introduced in the U.K. in 1966, making it the second country after the U.S. to adopt credit cards. The share of adults with a credit card is 60% in the UK (UK Cards Association, 2017) relative to 9.5% in Mexico (BBVA, 2015).

that the degree of misallocation should be decreasing in the financial stakes of the repayment decision. 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), the size of the repayment amount (which can be as high as £800 in a month), as well as a combined measure of the “financial stakes” of the repayment decision. We also show that the degree of misallocation is invariant to the time since account opening, suggesting that a learning model cannot explain observed behavior.⁷

The third part of the paper considers a number of different heuristics that might better explain repayment behavior. We consider a balance-matching heuristic under which the share of repayments on each card is matched to the share of balances on each card. Balance matching could arise from the salient placement of balances on credit card statements and the broad tendency for humans, and other species, to engage in “matching behavior” in related choice environments. Balance matching could also arise from individuals repaying a constant percentage of the balance on each card in a given month, a rule-of-thumb that would lead to inefficient behavior on both the allocative and extensive margins. We also consider five alternative heuristics, including a “1/N” heuristic (Benartzi and Thaler, 2001), in which payments are split evenly across cards, and the “debt snowball method,” in which payments are concentrated on the card with the lowest balance, as recommended by some financial advisors.⁸

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

⁷ Misallocation is similarly invariant to the number of days between payment dates, suggesting that frictions in coordinating repayments across cards with different due dates cannot explain repayment behavior.

⁸ For example, see syndicated radio host Dave Ramsey: <https://www.daveramsey.com/blog/get-out-of-debt-with-the-debt-snowball-plan>. Also see Amar et al. (2011) and Brown and Lahey (2015) for experimental studies that investigate this hypothesis.

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.⁹ In sensitivity analysis, we show that our results are robust to two potential threats to validity. First, we show that our results are not explained by anchoring on minimum payments (Keys and Wang, 2018; Navarro-Martinez et al., 2011; Stewart, 2009). Second, we show that our results are not driven by individuals using automatic payments (“autopay”), which is rarely used in our baseline sample.

We also evaluate each of our models in “horse race” type analysis where we determine the best fit model on an individual \times 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. In specifications where we allow for multiple types, balance matching is the best fit model for roughly half of the observations in our sample. 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.

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., proportional increase in R^2) in our machine learning models. Consistent with the balance matching results, we find that balances have the highest variable importance, with variable importance 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.

Our findings are related to a number of strands of research. The finding of non-optimal repayments is closely related to the aforementioned Ponce et al. (2017) study on credit card

⁹ The one exception is the $1/N$ heuristic, which captures some behavior exactly, and has a comparable goodness-of-fit with the root mean square error metric.

borrowing in Mexico, and more broadly related to a larger literature on non-optimal behavior in consumer financial markets (e.g., Agarwal et al., 2009; Keys et al., 2016; Jorring, 2018).¹⁰

Our investigation of optimal inattention models contributes to active literature on the explanatory power of these frameworks. For instance, consistent with our findings, Chetty et al. (2014) show that individuals in Denmark are inattentive in their pension contributions regardless of the economic stakes, pushing against models of optimal inattention. In contrast, in online experiments, Taubinsky et al. (2018) find sharp evidence that attention to taxes is increasing in tax rates, supporting these models. While we do not believe there is a single correct model for behavior, our results suggest that in the credit card repayment domain, optimal inattention models have little explanatory power.

The balance matching result relates to a well-established literature in psychology on matching behavior. In a seminal paper, Herrnstein (1961) showed that pigeons peck keys for food in proportion to the time it takes for the keys to rearm rather than concentrating their effort on the key that rearms most quickly. This behavior, which became known as the “matching law,” has since been documented across a number of species. More recently, a large experimental literature has shown evidence of “probability matching” in repeated choice experiments. For instance, when faced with repeated choices between a gamble with a 70% chance of winning and one with a 30% chance of winning, subjects choose gambles in proportion to the chances of winning, even though choosing the gamble with the higher payoff is a dominant strategy (see Vulkan, 2000 for a review). We discuss this literature, and our relation to it, in more depth in Section 5. To the best of our knowledge, our paper provides the first evidence of matching in the field.¹¹

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

¹⁰ Our paper is more broadly related to a literature that examines how credit card utilization responds to changes in contract terms and shocks to income (e.g., Gross and Souleles, 2002; Olafsson and Pagel, 2018; Agarwal et al., 2018).

¹¹ See DellaVigna (2009) for a review of the evidence on choice heuristics using field data.

loans) or make adjustments on the extensive margin (e.g., by adjusting the tradeoff between debt repayment and consumption). As we mention above, 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 (repaying a fixed percentage of the balances).

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 tests whether optimal inattention and learning can fit the observed behavior. Section 5 lays out alternative heuristics for debt repayment, including the balance-matching heuristic. Section 6 tests between these repayment models. Section 7 discusses the implications of these results for the industrial organization of credit card markets and public policy. 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 (see Footnote 1).¹²

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

¹² While we do not have data that would allow us to comprehensively examine the representativeness of the CCPS data, the distribution of interest rates in the CCPS is very similar to the aggregate distribution reported in UK Cards (2013), suggesting that the CCPS is roughly representative of the entire market.

credit card holders have two or more cards, and individuals with two or more cards account for 72.2% of outstanding balances (FCA, 2016).¹³

Our unit of analysis is the individual \times month. In the remainder of the paper, we refer to individual \times 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.¹⁴ 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.

We implement our sample restrictions at the individual \times month level, excluding months that do not satisfy our restrictions, but otherwise keeping these individuals in our sample. To focus on repayment decisions, we drop observations where individuals are delinquent or have defaulted on at least one card in their portfolio, or where individuals pay less than the minimum due or more than the full balance on at least one card.¹⁵ 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

¹³ 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.

¹⁴ In the two-card sample, two-thirds of observations have payment dates that are 10 days apart or fewer. See Figure A4 for a histogram of the difference in due dates between cards.

¹⁵ 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.

the manner described in Section 3, minimizes interest charges.^{16,17} These restrictions remove a further 25% 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 the third set of 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 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.¹⁸ 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

¹⁶ 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).

¹⁷ 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.

¹⁸ 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.

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

In this section, we compare 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.¹⁹

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.²⁰ We also do not focus on the optimality of revolving balance allocations because revolving balances are a “stock” that cannot typically be quickly adjusted.²¹ Thus, to determine whether 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

¹⁹ 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.

²⁰ 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.

²¹ 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.

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 (i) 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 Appendix 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.²² Alternatively put, over 90% of individuals should put 100% of their excess payments on the high interest rate card but fewer than 10% do so. Summary data for payments in excess of minimum are shown in Appendix 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 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.

As we mention in the introduction, these findings on misallocation are strikingly similar to those documented by Ponce et al. (2017) using a sample of Mexican credit card holders. In

²² 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.

particular, the most comparable result is that the share of payments in excess of the minimum misallocated to the high APR card is 50% among Mexican credit card holders and 46% among UK credit card holders. This remarkable similarity shows that misallocation among Mexican credit card holders is not due to their relatively limited experience with credit cards; in the UK market penetration is 66% and cards have been available since 1966.

3.1 Costs of Misallocation

What are the costs of the failure to optimize? Our baseline approach is to 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, we construct the optimal allocation 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, we construct the optimal allocation by allocating as much of the aggregate balance as possible to the credit card with the lowest interest rate, then allocating any remaining balance to the card with the next lowest rate, and so on. The interest savings from this exercise can be thought of as the interest savings that could be achieved if the individual was immediately and costlessly able to conduct the optimal “balance transfer” of balances across cards.

Panel B of Table 2 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.

In Appendix B, we present two additional sets of results on the cost of misallocation, focusing on the two card sample to keep the analysis trackable. First, we repeat the interest savings calculations above for cards that were excluded by our sample restrictions. For some excluded observations, such as those with no revolving balances, the interest savings are zero. For others, such as observations where individuals pay the minimum on both cards,

the counterfactual steady state yields interest savings. Total savings, which combine positive savings for borrowers and zero savings for non-borrowers, are roughly one-third those in the baseline sample.

Second, we consider an alternative approach to calculating the cost of misallocation where counterfactual interest costs are determined by simulating forward outcomes when individuals optimally repay their credit cards over time. This simulation is supposed to measure the gains from “learning” the optimal repayment rule, which the individual can then implement over time and incrementally shift their balances across cards. While interest savings are small at first, we show that over a 12 month time horizon, the interest savings from this incremental approach is quite similar to the interest savings from our baseline steady state counterfactual.

There are a number of caveats to these calculations. 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.

4 Optimization Frictions

In this section, we examine whether optimization frictions can explain repayment behavior. We start by examining the evidence for optimal inattention models and then turn to the evidence for models of learning.

4.1 Optimal Inattention

In optimal inattention models, individuals face a fixed cost of optimization and only “optimize” if the benefits from doing so are greater than this fixed cost (Sims, 2003). Specifically, in our context, individuals are already making positive credit card repayments on all of their 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.²³ 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 three measures of the economic stakes: (i) the difference in APR across cards, (ii) the total repayment made that month, and (iii) the “financial stakes” of the repayment decision, which is defined as the product of the difference in APR across cards and the total repayment made that month – and thereby captures the annualized financial stakes from the repayment decision. Since the gains from optimizing are increasing in each of these measures, under the fixed cost explanation, the percentage of misallocated repayments should be declining in these three measures. Moreover, for individuals with large economic stakes, we would expect the degree of misallocation to be close to zero.²⁴

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.²⁵ The flat relationship indicates that individuals are not less likely to misallocate repayments even when there is a large APR difference (more than 15 percentage points). 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. Panel C shows there is no evidence of a decreasing relationship between the percentage of misallocated payments and the financial stakes. The

²³ This contrasts with optimal mortgage refinancing (e.g., Keys et al., 2016), a setting where optimization would require actively soliciting a new mortgage offer.

²⁴ These measures of economic stakes are not randomly assigned. However, it is important to note that the lack of random assignment is only a threat to “external validity” and not to “identification”. Specifically, the fact that there are some individuals who do not optimize despite large incentives to do so implies that there are some individuals whose behavior cannot be rationalized by an inattention model. What requires a stronger assumption is to infer from this pattern that even individuals with small gains from optimization would still not optimize if they faced large incentives to do so.

²⁵ 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. See Chetty et al. (2014) for more details on the binned-scatter plot methodology.

relationship is slightly increasing for the same reason as the relationship with total payments. Table 3 presents summary statistics on actual, optimal and misallocated payments by quintiles of these measures of economics stakes.²⁶

4.2 Learning

An alternative 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 A of Figure 3 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.²⁷ This finding suggests that neither learning nor time-varying adjustment costs can explain the observed non-optimizing behavior.

A related explanation is that individuals have difficulty optimizing repayments across cards when the due dates are far apart in time. From a cognitive perspective, individuals may be less likely to process credit card repayments as a joint decision when the due dates are far apart. From a budgetary perspective, individuals may 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. Panel B of Figure 3 examines this explanation by showing a binned-scatter plot of the percentage of misallocated payments against the difference in

²⁶ Appendix Figure A2 and Appendix Table A4 present the analogous results for misallocated payments in excess of the minimum payment.

²⁷ Table 3 shows that summary statistics that correspond to this figure. Appendix Figure A3 and Appendix Table A4 present the analogous results for misallocated payments in excess of the minimum payment.

payment due dates.²⁸ The flat relationship indicates that these type of within-month allocation frictions cannot explain our result.

While it would be tempting to go beyond this simple analysis and try to quantify the fixed costs required to rationalize observed repayment behavior, doing so would require making strong, untestable assumptions. In particular, a key decision is whether to model optimization frictions as a fixed cost that needs to be paid every period (e.g., a transaction cost of adjusting repayments) or as a fixed cost that only needs to be paid once (e.g., a one-time cognitive cost of learning the optimal repayment rule). The estimated fixed costs would largely be determined by this modeling decision. If we assumed there was a period-by-period fixed cost, repayment behavior could be rationalized by modest optimization frictions. If, instead, we assumed there was a one-time fixed cost, the benefits of learning the optimal rule could be as large as the discounted sum of interest savings from optimal credit card payments over the life cycle, which would imply a very large fixed cost of adjustment. Since that data do not provide us with guidance on the appropriate modeling assumption, we are hesitant to undertake this exercise.

5 Heuristics

If individuals do not allocate their payments in an optimal – or constrained optimal – manner, 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 $k = 1 \dots K$ index cards, p_k indicate payments, and q_k indicate balances. Balance-matching payments are given by

$$\frac{p_k}{\sum_{\kappa=1}^K p_{\kappa}} = \frac{q_k}{\sum_{\kappa=1}^K q_{\kappa}} \quad \text{for } k = 1 \dots K \quad (1)$$

²⁸ Table 3 shows the summary statistics that correspond to this figure. Appendix Figure A3 and Appendix Table A4 present the analogous results for misallocated payments in excess of the minimum payment.

subject to the constraint that the individual pays at least the minimum and no more than the full balance on any of their cards.²⁹

Repayments consistent with balance matching could arise from the salient placement of balances on credit card statements and the broad tendency for human (and other species) to engage in “matching behavior” in related choice environments. As shown in Appendix Figure A5, balances are perhaps the most 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 saliently displayed balances (instead of less saliently displayed interest rates).^{30,31}

The literature on matching behavior traces back to Herrnstein (1961), who showed 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. Follow-up studies replicated Herrnstein’s “matching law” among pigeons (Miller, 1976; Baum, 1974) and documented similar behavior among rats (Jacquet, 1972), fish (Behrend and Bitterman, 1961), and humans when undertaking more complex tasks (e.g., Skinner et al., 1996). Writing at the time, Arrow (1958) noted that these results challenged conventional economic models of behavior.³²

Matching behavior has perhaps been most widely documented among humans in repeated lottery choice experiments. In a typical study, subjects are asked to make repeated choices between, say, a gamble with a 70% probability of winning and one with a 30% probability of winning. In behavior known as “probability matching,” subjects consistently place bets in proportion to the probabilities of winning, even though betting on the option with the highest probability first-order stochastically dominates any other decision. The most well-known

²⁹ 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.

³⁰ 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 (%).

³¹ 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 Appendix C and show that this channel does not explain our results.

³² “We have here an experimental situation which is essentially of an economic nature in the sense of seeking to achieve a maximum of expected reward, and yet the individual does not in fact, at any point, even in a limit, reach the optimal behavior. I suggest that this result points out strongly the importance of learning theory, not only in the greater understanding of the dynamics of economic behavior, but even in suggesting that equilibria maybe be different from those that we have predicted in our usual theory.” Arrow (1958) p.14.

papers in the literature are Rubinstein (2002), Gal (1996) and Loomes (1998). See Vulkan (2000) for a review paper that documents probability matching, and its limitations, across 22 repeated lottery experiments.

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 use balance matching as an “anchor” for their repayment behavior (Tversky and Kahneman, 1974). 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.³³

Finally, as we mention in the introduction, repayments consistent with balance matching could arise from individuals repaying a constant percentage of the balance on each card in a given month. If this percentage changes within individuals over time (e.g., 10% in January, 15% in February, etc.) it would be difficult to distinguish a constant percentage model from a balance matching model. If this percentage is fixed within individuals over time, we could distinguish between these models using random time-series variation in balances. Unfortunately, we do not have random time-series variation in balances, and the observed variation is likely to be driven by factors that also affect repayment decisions.

5.2 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.³⁴

³³ Ponce et al. (2017) show that there is a positive correlation between the share of payments towards cards and the share of debt held on cards, suggesting an approximate balance-matching rule may be at work among their sample of Mexican cardholders.

³⁴ By excess mass, we are using the terminology that Saez (2010) and Chetty et al. (2011) use to describe bunching in response to kinked tax schedules. In our context, the excess mass is the area in excess of what is predicted by a linear extrapolation between points on the histogram just to the left and the right of the 1/N values.

While it would be nice to conclude that 11.7% of individuals repay according to a $1/N$ rule, there are a number of factors which suggest a more nuanced interpretation of the data. First, the $1/N$ rule could be interpreted as an anchor that individuals do not implement exactly. For example, an individual who pays £250 on card A and £225 on card B might be intending to split repayments. To account for this possibility, we will measure the performance of the $1/N$ heuristic in our goodness of fit analysis below. This analysis will gauge how well the $1/N$ heuristic fits the data, even when it does not fit behavior exactly.

A second nuance arises from the fact that $1/N$ behavior is disproportionately concentrated among individuals who “round” their 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 4 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.³⁵

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%. Thus, at the extreme, 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 other extreme that all of these individuals would have allocated 50% on the high APR card if they

³⁵ The propensity to make round number payments is highly correlated across cards within an individual \times 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.

had counterfactually not rounded their payments.

Thus, the quantitative importance of $1/N$ depends upon what assumptions we are willing to make. If we think $1/N$ is something that individuals implement exactly, then the $1/N$ rule captures at most 11.7% of repayments, and possibly a lower percentage if we believe that some of this is due to rounding. If we view $1/N$ as an anchor that individuals approximate, then it is best to judge the empirical relevance of this measure in the goodness of fit analysis that follows.

5.3 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 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 the 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.³⁶ If an individual fully pays off a card, this heuristic has the additional benefit of “simplifying” the individual’s debt portfolio. See also Amar et al. (2011) and Brown and Lahey (2015) for laboratory evidence on this approach to debt repayment.

5.4 Interest Payments Under Different Heuristics

How costly would it be for consumers to follow these heuristics? We calculate annualized interest savings for the steady state balances that would arise from each of the different heuristic models. For example, if consumers follow a balance matching heuristic, then the process of making larger payments on the card with the larger balance would cause balances to asymptote to a steady state with equal balances on each card.³⁷

Appendix Table A5 shows the distribution of annualized interest savings (compared to actual interest payments) under each heuristic for the two-card sample. We show the interest saving from optimal payment on the first row as a benchmark and omit the $1/N$ heuristic because steady state balances are unchanged and interest savings are therefore zero. All the heuristics have median savings close to zero, but have wide dispersion in savings or losses

³⁶ For example, syndicated radio host Dave Ramsey argues: “But when you ditch the small debt first, you see progress. That one debt is out of your life forever. Soon the second debt will follow, and then the next. These little wins will give you a confidence boost, you’ll see that the plan is working, and you’ll stick to it” (see <https://www.daveramsey.com/blog/top-5-debt-snowball-questions-answered>).

³⁷ The steady state outcomes of the other heuristics are as follows. Under the $1/N$ heuristic, payments are equal across cards, and the steady state balances are unchanged relative to the status quo. Under Heuristic 1 (repay the card with the lowest capacity), payments are prioritized to the card with the lowest capacity, and in steady state capacities on the two cards are equalized, subject to credit limits. Under Heuristic 2 (repay the card with the highest capacity), payments are prioritized to the card with the highest capacity, and in steady state balance are reduced as much as possible on that card, subject to credit limits. Under Heuristic 3 (repay the card with the highest balance), payments are prioritized to the card with the highest balance, and in steady state the balances on the two cards are equalized, subject to credit limits. Under Heuristic 4 (repay the card with the lowest balance), payments are prioritized to the card with the lowest balance, and in steady state balance are minimized on that card, subject to credit limits.

across observations. This wide dispersion occurs because the heuristics prioritize payments to cards based on balances (and credit limits), which is not systematically correlated with the optimal behavior of prioritizing payments toward the high APR card.

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). The goal of this analysis is to find the single model that best fits the data. Second, we evaluate the performance of our models in “horse race” type analysis where we determine the best fit model on an observation-by-observation basis. We conduct binary comparisons where we compare each model to a single alternative, and horse races with 3 or more models, which allows us to classify observations into heterogenous types. We also examine the persistence of the best fit model within individuals over time.

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 and the bottom row focuses on the sample with non-round number payment amounts. The sample with round number payment amounts is shown in Figure A6.

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%. 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 stronger in the non-round payment sample (and weaker in the round number payment sample).³⁸

³⁸ 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. While the heaping in Panel A of Figure 4 provides clear evidence of rounding, estimating a two-stage

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).³⁹ 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.⁴⁰ 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 A.

model where individuals round in the second stage is difficult because the second stage rounding function is not “invertible.” That is, it is hard to know whether an individual who is paying £100 would have counterfactually paid £95 or £105 if they had not rounded their payment.

³⁹ Pearson’s ρ is also the square root of the R^2 from a univariate regression of actual payments on predicted payments.

⁴⁰ 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.

Figure 7 shows these measures of goodness-of-fit under the different models. (Appendix Table A6 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 optimal model does generate a 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.⁴¹

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 A7 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 in Footnote 38.

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 easy to understand, has a psychological underpinning based on 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.

⁴¹ Appendix Figure A8 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.

6.2 Horse Races Between Alternative Models

We next evaluate these models using “horse races” where we determine the best fit model on an observation-by-observation basis. A model that fits a smaller number of observations very poorly, but a larger number quite well, might perform poorly under the goodness-of-fit analysis, but would perform well under this approach.

Our ability to identify the best-fit model on an observation-by-observation basis 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 constitute a significant empirical challenge.

Table 4 shows results of this horse race analysis in the pooled sample of individual \times 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 perform only slightly better than the uniform benchmark. The machine learning models have the best fit for between 69.0% and 76.5% of observations, which is similar to balance matching.⁴²

Panel B of Table 4 compares each of the models one-by-one to the balance-matching model.⁴³ 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 40.9% to 50.0% of observations.⁴⁴

⁴² 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.

⁴³ We exclude a comparison of the balance-matching and the uniform model, since it was shown in Panel A.

⁴⁴ In unreported analysis, we examined whether there was correlation between the best fit model and postal

The binary comparisons are useful for finding the single model that fits the largest percentage of observations. A complementary approach is to conduct horse races with 3 or more models, which allows us to classify observations into heterogeneous types. One challenge with this exercise is that – simply by chance – virtually any model will have the best fit for some observations. To be able to interpret the results as reflecting heterogeneity in types, we need to be disciplined in choosing which models to include and to also compare the percentage of observations fit by a given model to the uniform benchmark.

Appendix Table A8 shows the results of this exercise. In Column 1, we compare the optimal, $1/N$, and balance matching models. These are the models that seem plausible based on economic theory (optimal) or evidence ($1/N$ and balance matching). In column 2, we add the uniform benchmark to the set of models. In both specifications, balance matching has the best fit for half of the observations. Optimal and $1/N$ capture slightly more than one-fifth of observations when we do not include the uniform benchmark, but perform only marginally better than the uniform benchmark when it is included.

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 5 shows the within-person transition matrix for the best-fit model. The sample is restricted to individual \times 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 individuals 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, 23.2% 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,

code characteristics (such as postal code income and education). We were unable to detect any economically or statistically significant correlations. However, we caution against drawing too strong conclusions about heterogeneity, because we do not know whether postal code characteristics are good proxies for the demographic characteristics of our sample of individuals.

which is fixed to be constant within an individual over time, 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% 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 $1/N$ repayments (or the tendency to round payments) is a stable feature of individual behavior.

6.3 Sensitivity Analysis

In Appendix C, we present two sets of sensitivity analysis. First, for some credit cards, minimum payments are set as a proportion of balances. For these cards, balance matching behavior could instead arise from individuals making payments that are proportional to minimum payments. We rule out this explanation by showing that the balancing matching result holds among cards where minimum payments are not proportional to balances. Second, we investigate the sensitivity of our results to whether the individual uses automatic payment (“autopay”). We show that autopay is rare in our baseline sample (11%). 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.

6.4 Balances and APRs in Machine Learning Models

In addition to providing us with an upper benchmark, the machine learning models allow us to assess the relative importance of balances and APRs in predicting repayment behavior. Specifically we calculate the variable importance, which can be thought of as the incremental increase in R^2 from adding a given variable to the model. Appendix Table A9 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. Appendix Table A10 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 A11 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 from Appendix Table A10, indicating that our variable importance measures are not particularly sensitive to random variation in the training dataset.

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 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 Discussion

The fact that individuals do not repay credit cards optimally, and instead repay in proportion to balances, has implications for the industrial organization of credit card markets and public policy. While it is beyond the scope of the paper to rigorously examine these implications, we provide some initial thoughts below.

To think through the implications for consumer and producer surplus, consider a two-stage model of credit card market. In stage 1, individuals shop for and select credit cards. In stage 2, individuals spend on and repay their credit cards. The analysis in this paper implies that conditional on reaching stage 2, individuals do not optimally concentrate their borrowing on their cheapest card. From the perspective of the issuer, this implies that being the high cost card in a consumers' portfolio can generate large stage 2 profits. However, to understand the full

effect of non-optimal repayments on producer and consumer surplus, we need to understand how these stage 2 profits affect behavior in stage 1.

One possibility is that firms dissipate their stage 2 profits via competition for an “install base” of consumers in stage 1. If stage 2 profits lead firms to set very low stage 1 prices (e.g., teaser rates) to acquire an install base, then it is theoretically possible for firms to dissipate all of their stage 2 rents, and for consumer surplus to be similar to what it would be in a counterfactual scenario where consumers repay their balances optimally (although the timing of borrowing will still be distorted from the teaser rate). A second possibility is that firms dissipate some – perhaps most – of their stage 2 profits in acquisition behavior. The large number of mail solicitations, aggressive new card promotions, and heavy television advertising in the industry is consistent with this theory. To the extent that much of these acquisition costs are businesses stealing from rivals, and not category expansion, this expenditure can be thought of as a deadweight loss. A third possibility is that firms do not dissipate much of their stage 2 profits in stage 1. This could occur because there are decreasing returns to acquisition activity (conditional on sending a mail solicitation every week, the marginal return to an additional mailer might be quite small). Adjudicating between these different possibilities is important but beyond the scope of this paper.

From a policy perspective, a natural response to the observed non-optimality of repayments is to make interest rates more salient. While Seira et al. (2017) found no effect of interest rate salience on indebtedness and defaults in their study of Mexican credit card holders, there may be alternative ways of presenting interest rate information that successfully shift behavior. A broader policy response is to make it easier for consumers to share credit card data with third parties, such as FinTech firms.⁴⁵ In principle, these firms could use these data to help consumers optimize their repayments across their credit cards, although there is no guarantee that they will nudge consumers in an optimal direction.

⁴⁵ For instance, recently implemented Open Banking regulation in the UK requires large banks, at their consumers’ discretion, to give third parties access to account data through secure application programming interfaces (APIs).

8 Conclusion

In this paper, we used linked data on multiple cards from five major credit card issuers to study how individuals repay their debt. First, 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%. These findings are almost identical to Ponce et al. (2017), and indicate that the behavior they documented did not merely reflect limited consumer experience with credit cards in the Mexican market.

Second, we showed repayment behavior was inconsistent with models of optimization frictions. Specifically, we showed that degree of misallocation was invariant to the financial stakes of the repayment decision, implying that models of optimal inattention with a fixed cost of adjustment could not explain our findings. We also showed that the degree of misallocation is invariant to the time since account opening, suggesting that a learning model cannot explain observed behavior.

The third part of the paper showed that repayment behavior is consistent with a balance-matching heuristic under which the share of repayments on each card is matched 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. Making repayments in proportion to balances could arise from broad tendency for humans, and other species, to engage in “matching behavior” in related choice environments, but could also arise from individuals repaying a constant percentage of their balances across cards. Unpacking the underlying explanations for repayment behavior is a fruitful area for future research.

References

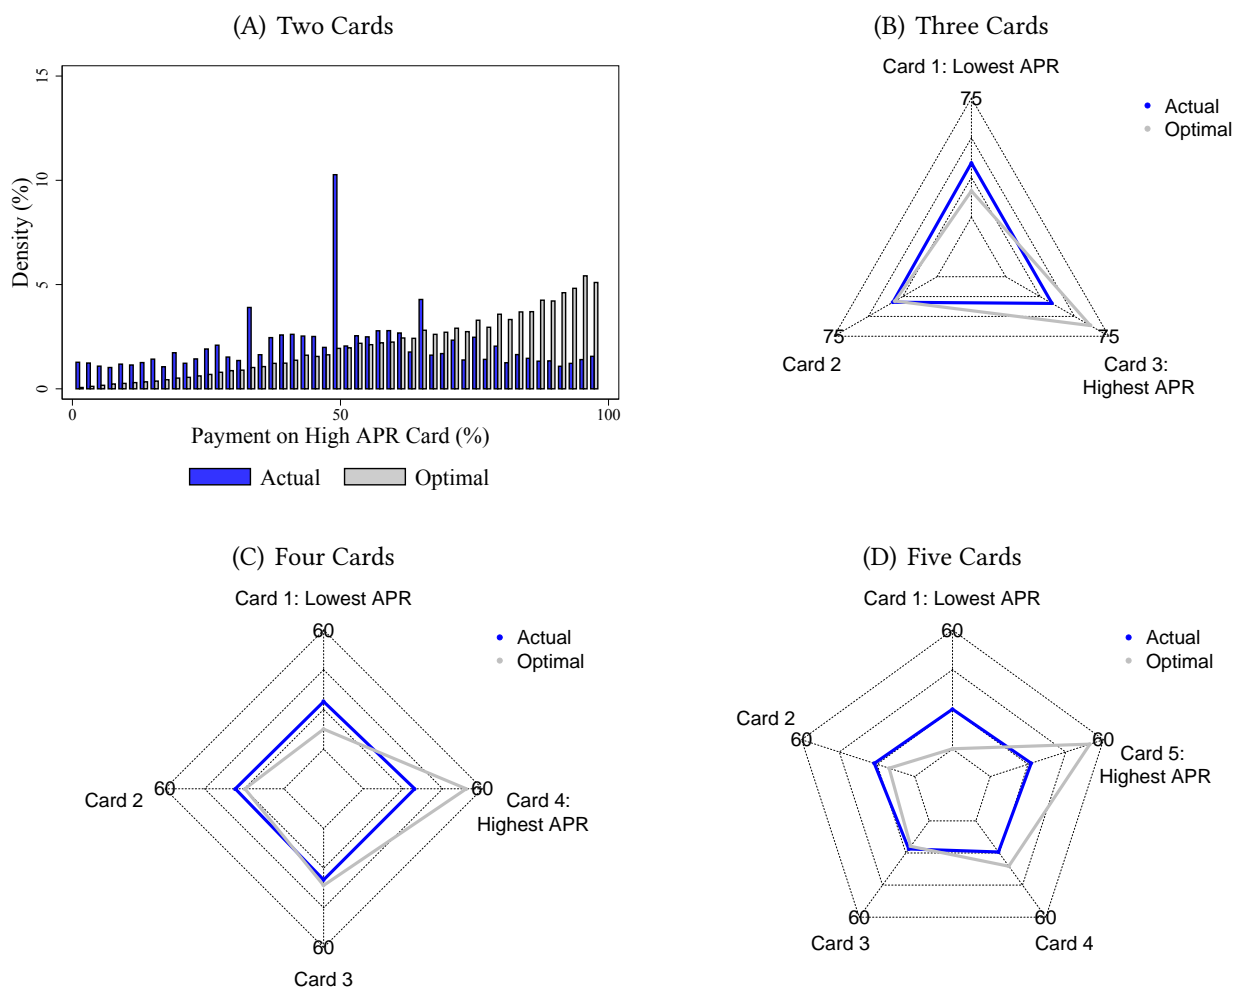
- Agarwal, S., S. Chomsisengphet, N. Mahoney, and J. Stroebel (2015). Regulating Consumer Financial Products: Evidence from Credit Cards. *The Quarterly Journal of Economics* 130(1), 111–164.
- Agarwal, S., S. Chomsisengphet, N. Mahoney, and J. Stroebel (2018). Do Banks Pass through Credit Expansions to Consumers Who want to Borrow? *The Quarterly Journal of Economics* 133(1), 129–190.

- Agarwal, S., P. M. Skiba, and J. Tobacman (2009). Payday Loans and Credit Cards: New Liquidity and Credit Scoring Puzzles? *American Economic Review* 99(2), 412–17.
- Amar, M., D. Ariely, S. Ayal, C. E. Cryder, and S. I. Rick (2011). Winning the Battle but Losing the War: The Psychology of Debt Management. *Journal of Marketing Research* 48(SPL), S38–S50.
- 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.
- Arrow, K. J. (1958). Utilities, Attitudes, Choices: A Review Note. *Econometrica* 26(1), 1–23.
- BBVA (2015). Mexico Banking Outlook. Technical report, BBVA Research.
- Behrend, E. and M. Bitterman (1961). Probability-Matching in the Fish. *The American Journal of Psychology* 74(4), 542–551.
- Benartzi, S. and R. H. Thaler (2001). Naive Diversification Strategies in Defined Contribution Saving Plans. *American Economic Review* 91(1), 79–98.
- Breiman, L. (2001). Random forests. *Machine Learning* 45(1), 5–32.
- Brown, A. L. and J. N. Lahey (2015). Small Victories: Creating Intrinsic Motivation in Task Completion and Debt Repayment. *Journal of Marketing Research* 52(6), 768–783.
- 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, S. Leth-Petersen, T. H. Nielsen, and T. Olsen (2014). Active vs. Passive Decisions and Crowd-Out in Retirement Savings Accounts: Evidence from Denmark. *The Quarterly Journal of Economics* 129(3), 1141–1219.
- Chetty, R., J. N. Friedman, T. Olsen, and L. Pistaferri (2011). Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records. *The Quarterly Journal of Economics* 126(2), 749–804.
- 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). Greedy function approximation: A gradient boosting machine. *The Annals of Statistics* 29(5), 1189–1232.
- Gal, I. (1996). Understanding Repeated Simple Choices. *Thinking & Reasoning* 2(1), 81–98.
- Gross, D. B. and N. S. Souleles (2002). Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data. *The Quarterly Journal of Economics* 117(1), 149–185.
- 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.
- Jacquet, Y. F. (1972). Schedule-Induced Licking During Multiple Schedules. *Journal of the Experimental Analysis of Behavior* 17(3), 413–423.
- Jorring, A. (2018). The Costs of Financial Mistakes: Evidence from US Consumers. *Working Paper*.
- 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 (2018). Minimum Payments and Debt Paydown in Consumer Credit Cards. *Journal of Financial Economics* (forthcoming).
- Loomes, G. (1998). Probabilities vs money: A test of some fundamental assumptions about rational decision making. *Economic Journal* 108(447), 477–489.
- Mullianathan, S. and J. Spiess (2017). Machine Learning: An Applied Econometric Approach. *Journal of Economic Perspectives* 31, 87–106.
- Navarro-Martinez, D., L. C. Salisbury, K. N. Lemon, N. Stewart, W. J. Matthews, and A. J. L. Harris (2011). Minimum required payment and supplemental information disclosure effects on consumer debt repayment decisions. *Journal of Marketing Research* 48, S60–S77.
- Olafsson, A. and M. Pagel (2018). The liquid hand-to-mouth: Evidence from personal finance management software. *The Review of Financial Studies*, hhy055.
- Ponce, A., E. Seira, and G. Zamarripa (2017). Borrowing on the Wrong Credit Card: Evidence from Mexico. *American Economic Review* 107(4), 1335–1361.

- Rubinstein, A. (2002). Irrational Diversification in Multiple Decision Problems. *European Economic Review* 46(8), 1369–1378.
- Saez, E. (2010). Do taxpayers bunch at kink points? *American Economic Journal: Economic Policy* 2(3), 180–212.
- Seira, E., A. Elizondo, and E. Laguna-Müggenburg (2017). Are Information Disclosures Effective? Evidence from the Credit Card Market. *American Economic Journal: Economic Policy* 9(1), 277–307.
- Sims, C. A. (2003). Implications of Rational Inattention. *Journal of Monetary Economics* 50(3), 665–690.
- Skinner, C. H., S. L. Robinson, G. A. Johns, P. Logan, and P. J. Belfiore (1996). Applying Herrnstein’s Matching Law to Influence Students’ Choice to Complete Difficult Academic Tasks. *The Journal of Experimental Education* 65(1), 5–17.
- Stewart, N. (2009). The cost of anchoring on credit card minimum payments. *Psychological Science* 20, 39–41.
- Taubinsky, D., A. Rees-Jones, et al. (2018). Attention Variation and Welfare: Theory and Evidence from a Tax Salience Experiment. *Review of Economic Studies* (forthcoming).
- Tversky, A. and D. Kahneman (1974). Judgment under Uncertainty: Heuristics and Biases. *Science* 185(4157), 1124–1131.
- UK Cards (2013). UK Plastic Cards 2013. Technical report, The UK Cards Association.
- UK Cards Association (2017). UK Card Payments. Technical report.
- Vulkan, N. (2000). An Economist’s Perspective on Probability Matching. *Journal of Economic Surveys* 14(1), 101–118.

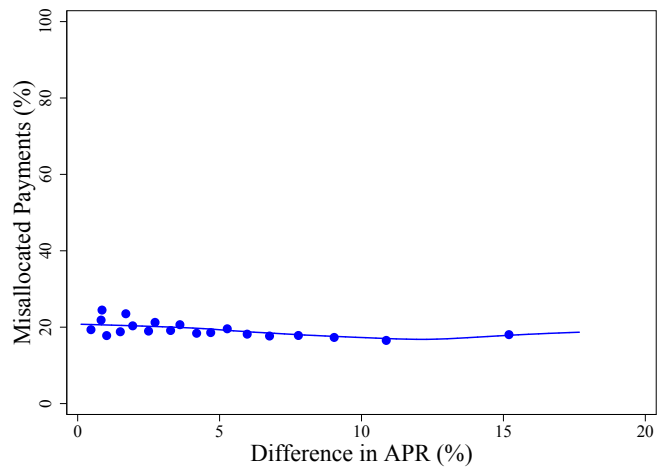
Figure 1: Actual and Optimal Payments




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 \times months in which individuals face an economically meaningful allocative decision. Sample sizes in Panel A: 394,111; Panel B: 264,420; Panel C: 83,880; Panel D: 19,120. See Section 2.2 for details on the sample construction.

Figure 2: Misallocated Payments by Economic Stakes


(A) Misallocated vs. Difference in APR



(B) Misallocated vs. Total Payments




(C) Misallocated vs. Δ APR \times Total Payment



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

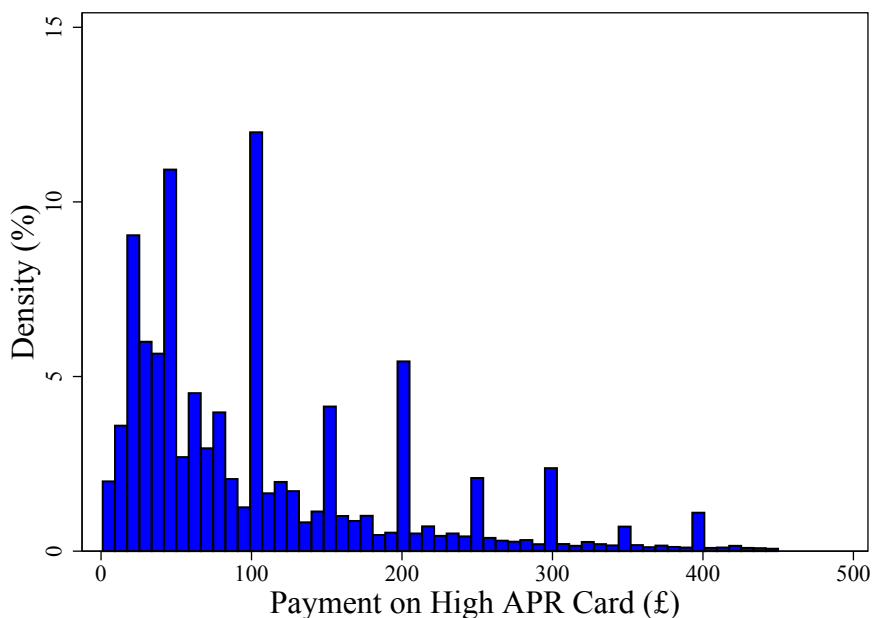
Figure 3: Misallocated Payments by Card Age and Difference in Due Dates




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

Figure 4: 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 35 for details. 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.

Figure 5: Balance Matching

(A) Baseline Sample



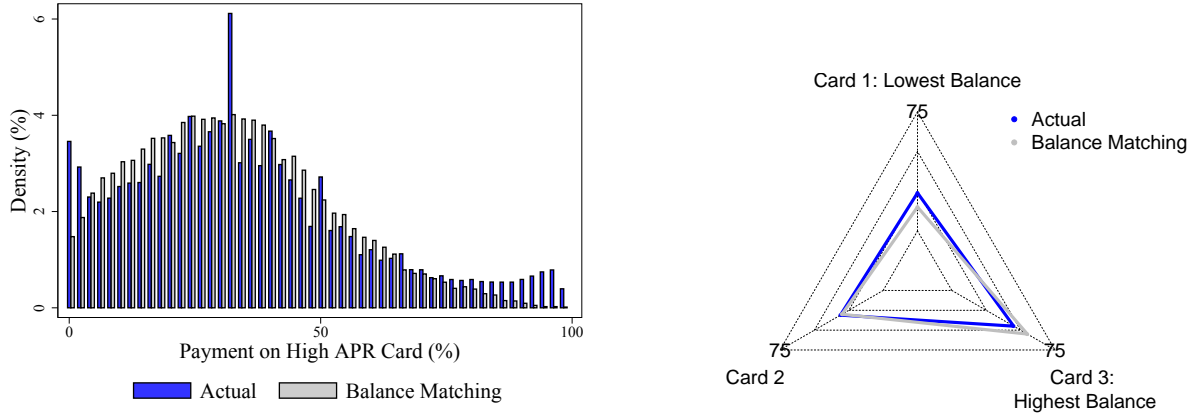
(B) Non-Round Number Payment Sample



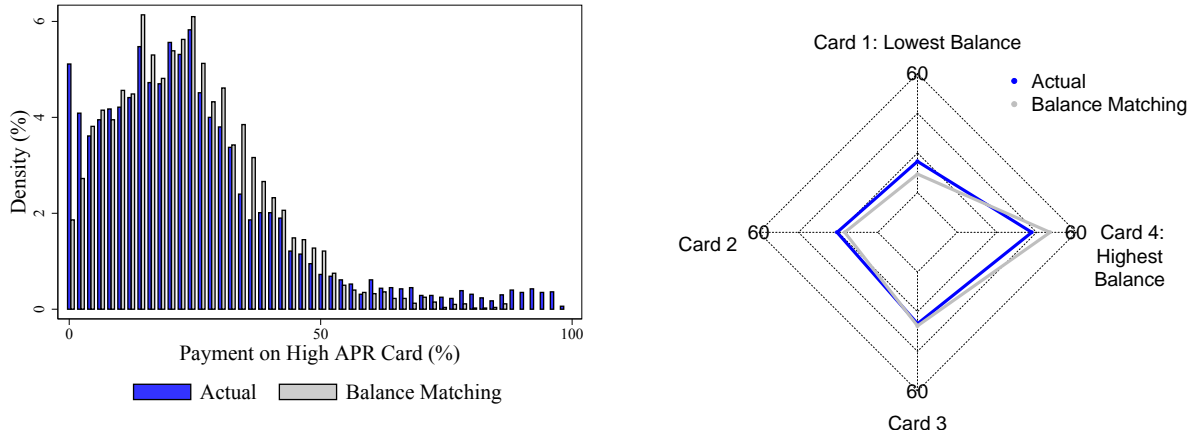
Note: Left panels show the distribution of actual and balance-matching payments on the high APR card. Right panels show the joint density of actual and balance-matching payments. Panel A shows the baseline sample, Panel B restricts the sample to non-round payment amounts (not multiples of £50). Round and non-round samples are defined by repayments on the high APR card. See Footnote 35 for details. The two-card sample is restricted to individual \times 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

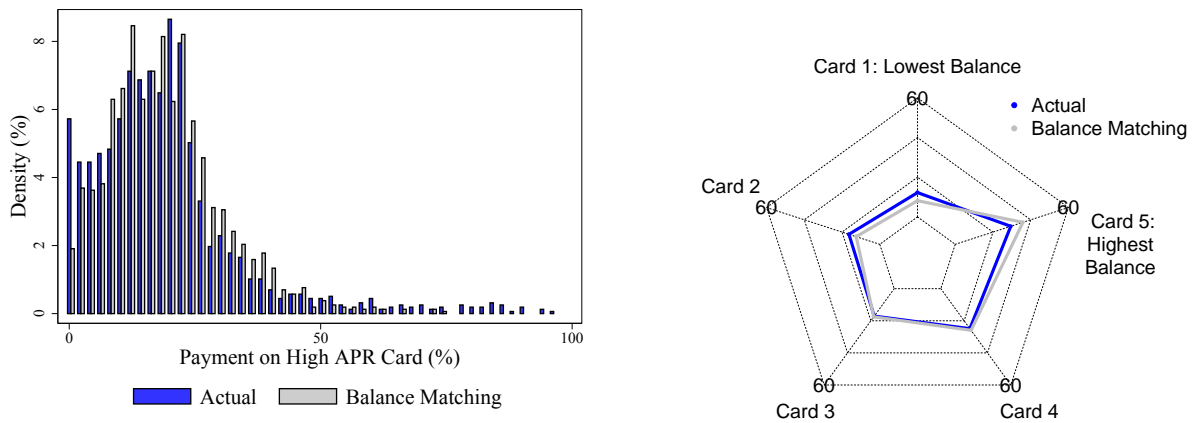
(A) Three Cards



(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 \times 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 \times 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.

Table 1: Summary Statistics

	(1) High APR Card		(2) Low APR Card		(3) Difference	
	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
<i>Card Characteristics</i>						
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
<i>Spending (£)</i>						
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
<i>Payments (£)</i>						
Repayments	259.76	733.92	229.69	657.60	30.07	913.65
<i>Interest Paid (£)</i>						
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
<i>Card Cycle (£)</i>						
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
<i>Card Status</i>						
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 \times 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 Payments and Optimal Payments, Interest Savings from Optimizing

Panel (A)

Actual and Optimal Payments on the High APR Card							
	Mean	Std. Dev.	Percentiles				
			10th	25th	50th	75th	90th
<i>i) As % Total Monthly Payment</i>							
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
<i>ii) Payment in £</i>							
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

Panel (B)

Annualized Interest Savings from Optimizing Credit Card Repayments				
	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: Panel A shows 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. $N = 394,111$. Percentiles and other summary statistics are taken with respect to individual \times month observations. Panel B shows summary statistics for annualized interest savings from a counterfactual “steady state” where individuals optimize balances across the credit cards we observe in our data, subject to the constraint of not exceeding their credit limits. The two-card sample has $N = 394,111$, the three-card sample has $N = 264,420$, the four card sample has $N = 83,880$, and the five-card sample has $N = 19,120$. In both panels, the sample is 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 $\text{£}1 = \text{\$}1.32$ at the midpoint of our sample period.

Table 3: Actual and Optimal Payments on High APR Card by Quintiles of Economic Stakes and Card Age

	Quintiles of $i - v$				
	1	2	3	4	5
<i>i) Difference in APR</i>					
Difference in APR (%)	0.81	2.03	3.78	6.44	14.98
Actual Payment (%)	49.79	50.27	50.98	51.86	53.19
Optimal Payment (%)	71.12	70.78	70.34	70.17	71.26
Difference (%)	21.33	20.51	19.36	18.32	18.07
<i>ii) Total Payment</i>					
Total Payment (£)	63.82	125.98	204.23	350.22	1658.91
Actual Payment (%)	49.70	50.09	51.04	52.62	52.64
Optimal Payment (%)	60.47	66.89	70.04	74.19	81.80
Difference (%)	10.78	16.80	19.00	21.56	29.16
<i>iii) Financial Stakes</i>					
Financial Stakes (£)	1.20	3.70	8.08	17.50	101.91
Actual Payment (%)	49.79	50.19	50.20	51.49	54.42
Optimal Payment (%)	64.67	68.13	68.92	72.20	79.75
Difference (%)	14.88	17.94	18.71	20.72	25.33
<i>iv) Difference in Due Dates</i>					
Difference in Due Dates (Days)	1.03	3.48	6.48	10.83	17.93
Actual Payment (%)	50.83	51.38	51.39	51.30	51.21
Optimal Payment (%)	70.07	70.71	70.65	70.99	71.20
Difference (%)	19.25	19.33	19.27	19.69	20.00
<i>v) Age of High APR Card</i>					
Age of High APR Card (Months)	3.82	5.57	7.50	9.50	11.48
Actual Payment (%)	43.58	43.54	43.65	43.31	43.96
Optimal Payment (%)	72.94	72.83	72.49	71.62	71.44
Difference (%)	29.36	29.29	28.84	28.31	27.48

Note: Summary statistics for actual and optimal payments on the high APR card by quintiles of economic stakes and card age. Cells report mean values within the quintile. The sample is restricted to individual \times months in which individuals face an economically meaningful allocative decision. $N = 394,111$. See Section 2.2 for more details on the sample construction. The exchange rate was $\text{£}1 = \$1.32$ at the midpoint of our sample period.

Table 4: Horse Races Between Alternative Models

Panel (A): Uniform vs. Other Rules										
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
<i>Win %</i>										
Uniform	32.87	46.66	45.53	50.32	45.44	47.48	47.09	31.01	23.55	27.14
Balance Matching	67.13									
1/N		53.34								
Optimal			54.47							
Heuristic 1 (Pay Down Lowest Capacity)				49.68						
Heuristic 2 (Pay Down Highest Capacity)					54.56					
Heuristic 3 (Pay Down Highest Balance)						52.52				
Heuristic 4 (Pay Down Lowest Balance)							52.91			
Decision Tree								68.99		
Random Forest									76.45	
XGB										72.86

Panel (B): Balance Matching vs. Other Rules									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
<i>Win %</i>									
Balance Matching	60.90	68.97	73.09	67.22	76.65	64.70	50.02	40.87	45.55
1/N	39.10								
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							49.98		
Random Forest								59.13	
XGB									54.45

Note: Table shows percentage of individual \times 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 \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.

Table 5: Transition Matrix for Best-Fit Model

	<i>Current Period</i>							
	Balance			H 1	H 2	H 3	H 4	1/N Rule
	Uniform	Matching	Optimal					
Uniform	23.18%	56.71%	0.24%	1.53%	0.35%	2.24%	0.59%	15.18%
Balance Matching	8.24%	83.06%	0.13%	1.37%	1.18%	1.07%	1.13%	3.82%
Optimal	9.09%	81.82%	9.09%	0.00%	0.00%	0.00%	0.00%	0.00%
Heuristic 1 (Pay Down Lowest Capacity)	7.98%	49.08%	0.00%	28.22%	0.00%	1.84%	4.91%	7.98%
Heuristic 2 (Pay Down Highest Capacity)	4.27%	66.67%	0.00%	0.85%	24.79%	1.71%	0.00%	1.71%
Heuristic 3 (Pay Down Highest Balance)	10.28%	47.66%	0.93%	4.67%	0.93%	17.76%	0.00%	17.76%
Heuristic 4 (Pay Down Lowest Balance)	11.59%	57.97%	0.00%	4.35%	0.00%	1.45%	20.29%	4.35%
1/N Rule	18.99%	27.89%	0.13%	2.39%	0.40%	1.33%	0.27%	48.61%

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

A Machine Learning Models

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

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

We use the `r` package “`rpart`” to fit the decision-tree model.⁴⁶ 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.⁴⁷ Appendix Figure A9 shows the estimated decision tree.

Random Forest The machine learning literature has proposed several variations on the tree model. One approach which has been found to work very well in practice is random forest

⁴⁶ See <https://cran.r-project.org/web/packages/rpart/vignettes/longintro.pdf> for a complete description of the function.

⁴⁷ 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.⁴⁸ 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.⁴⁹ 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.⁵⁰

⁴⁸ See <https://cran.r-project.org/web/packages/randomForest/randomForest.pdf> for a complete description of the function.

⁴⁹ See <http://cran.fhcrc.org/web/packages/xgboost/vignettes/xgboost.pdf> for a complete description of the function.

⁵⁰ 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.

B Costs of Misallocation: Extensions

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

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

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

An important complication with applying this alternative model is that we have to take a stand on counterfactual spending behavior. Assuming their spending “stays the same” is not an option: With the updated repayments, some individuals will go over their credit limits on the lower APR card. One option is to assume that individuals allocate their spending optimally. The optimal allocation of spending is achieved by prioritizing spending to the low APR card, allocating spending only to the high APR card once the credit limit of the low APR card is reached. If you told someone how to repay optimally, a good guess is that they would also

adjust their spending towards the optimal allocation (although we agree that they might not get all the way to optimal). Another option is to assume that they hold their spending fixed, unless they bump into a credit limit, in which case we can reallocate their spending to the other card. We show counterfactual interest payments under both of these assumptions for spending behavior.

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

Table A12 shows summary statistics for interest savings. The top panel shows savings when optimizing both payments and spending, the bottom panel shows savings from optimizing payments only. Interest savings in both versions of the dynamic optimal model are weakly positive. When individuals optimize both payments and spending (Panel A), mean annualized interest savings at 12 months are close to those from the steady state calculation (£58 versus £65) and the percentiles of the distribution are also similar. When individuals only optimize payments, balances take longer to converge to the steady state optimal allocation, and at a 12 months time horizon savings are one-third lower than savings when individuals optimize on both margins.

C Sensitivity Analysis

C.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 A5). If repayments are determined by a minimum-payment-matching heuristic, and minimum payments are proportional to balances, then mini-

imum payment matching could produce the observed repayment behavior.⁵¹

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 maximum of a fixed amount and a percentage of the balance. For instance, a typical minimum payment formula might be:

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

Consider the following scenarios for an individual with two cards:

- (i) If minimum payments are on the “fixed” part of the formula (balances greater than £1,250), and the percentages are identical (2% for both cards), then the balance-matching and minimum-payment-matching payments will be almost perfectly correlated.⁵²
- (ii) If the percentages differ, then balance-matching and minimum-payment-matching payments will be correlated, but to a lesser extent.
- (iii) If minimum payments are on the “percentage” part of the formula (balances less than £1,250), then the balance-matching allocation will not be correlated with the minimum-payment-matching allocation.

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

Figure A10 shows binned-scatter plots of actual and predicted payments on the high

⁵¹ Setting payments at multiples of the minimum amount (e.g., twice the minimum on each card) would also produce the observed repayment behavior.

⁵² 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.

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

In the same percentage sample, the balance-matching and the minimum-payment-matching payments are near-perfectly correlated ($\rho = 0.96$). As a result, the correlation between actual and balance-matching payments ($\rho = 0.63$) is nearly identical to the correlation between actual and minimum-payment-matching payments ($\rho = 0.61$). In the different percentages sample, the balance-matching and the minimum-payment-matching payments are more weakly correlated ($\rho = 0.86$), and the correlation between actual and balance-matching payments ($\rho = 0.41$) is stronger than the correlation between actual and minimum-payment-matching payments ($\rho = 0.28$). In the fixed sample, there is a much weaker correlation between the balance-matching payments and the minimum-payment-matching payments ($\rho = 0.56$), and the correlation between actual and balance-matching payments ($\rho = 0.50$) is substantially stronger than the correlation between actual and minimum-payment-matching payments ($\rho = 0.23$).

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

C.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 A11 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).⁵³ 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 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 (45.5% versus 45.7%). 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.⁵⁴ Summary statistics for actual and excess payments by autopay status are shown in Appendix Table A14.

Appendix Table A15 and Table A16 show our standard measures of model performance by whether individuals use autopay on both cards and do not use autopay on either card.⁵⁵ In particular, Appendix Table A15 shows our measures of goodness-of-fit (root mean square

⁵³ 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.


⁵⁴ Specifically, while balances are slightly higher in the autopay sample (£6,900 versus £5,800), repayments are substantially lower (£200 versus £510).

⁵⁵ Results are shown using the 20% hold-out sample.

error, mean absolute error, Pearson's correlation) for uniformly distributed repayments, optimal repayments, and balance-matching repayments separately for the autopay and non-autopay samples. Appendix Table A16 shows the results of horse-race analysis that compares uniformly distributed versus balance-matching payments, and balance-matching versus optimal repayments, separately for the autopay and non-autopay samples. While the exact results vary, the optimal model performs poorly and the balance-matching model performs well across all of these different measures of model performance in both the autopay and non-autopay sample. Thus, we conclude that our results are not particularly sensitive to whether individuals use autopay.

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

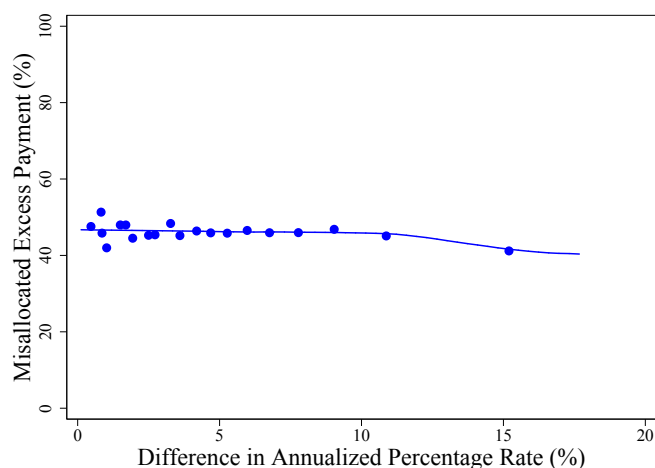
Figure A1: Actual and Optimal *Excess* Payments




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 \times months in which individuals face an economically meaningful allocative decision. See Section 2.2 for more details on the sample construction.

Figure A2: Misallocated *Excess* Payments by Economics Stakes

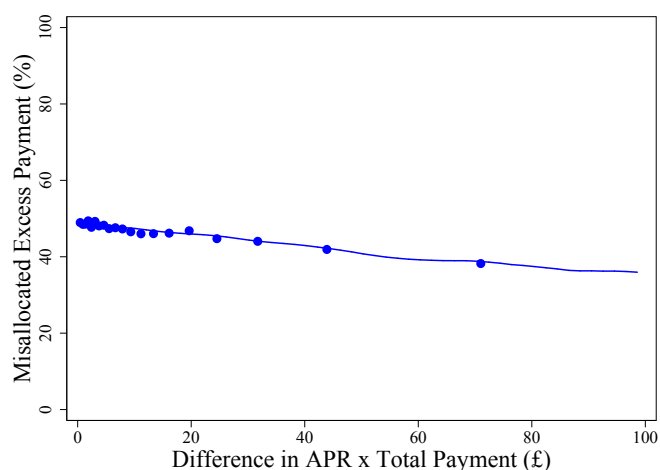
(A) Misallocated vs. Difference in APR



(B) Misallocated vs. Total Payments

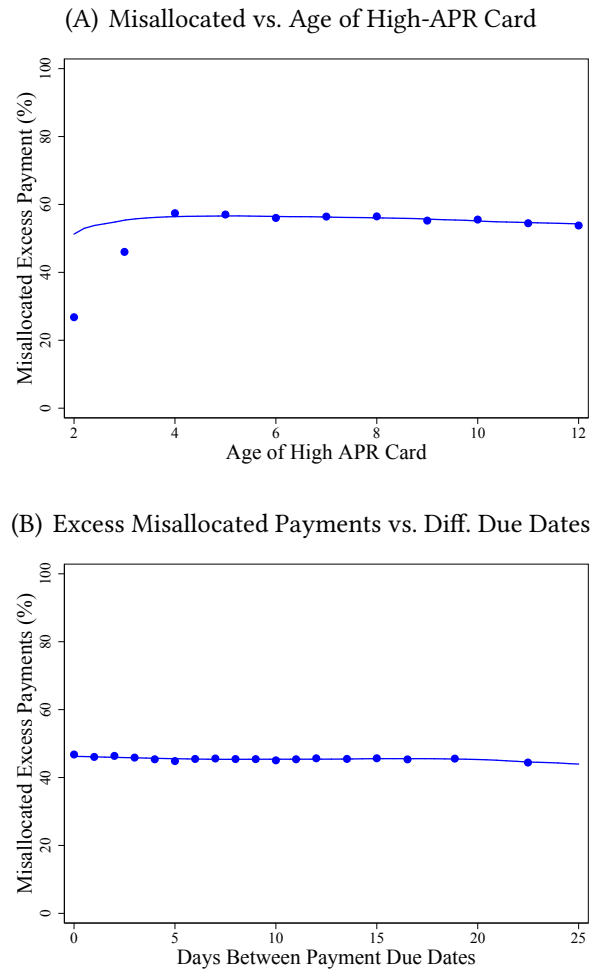


(C) Misallocated vs. Δ APR \times Total Payment




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

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



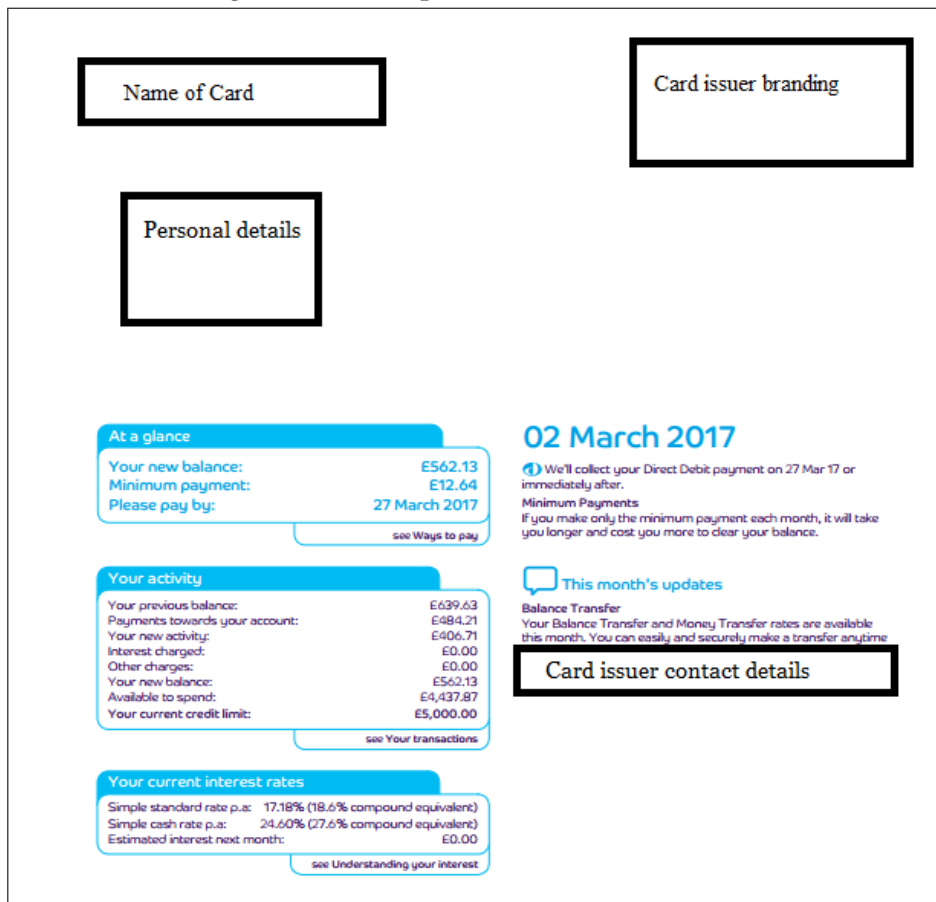
Note: Figure shows binned-scatter plots (with 20 equally sized bins) of misallocated payments in excess of the minimum payment against the difference in payment due dates (Plot A) and age of the high APR card (Plot B). Local polynomial lines of best fit, based on the non-binned data, are also shown.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

Figure A4: Histogram of Difference in Due Dates



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

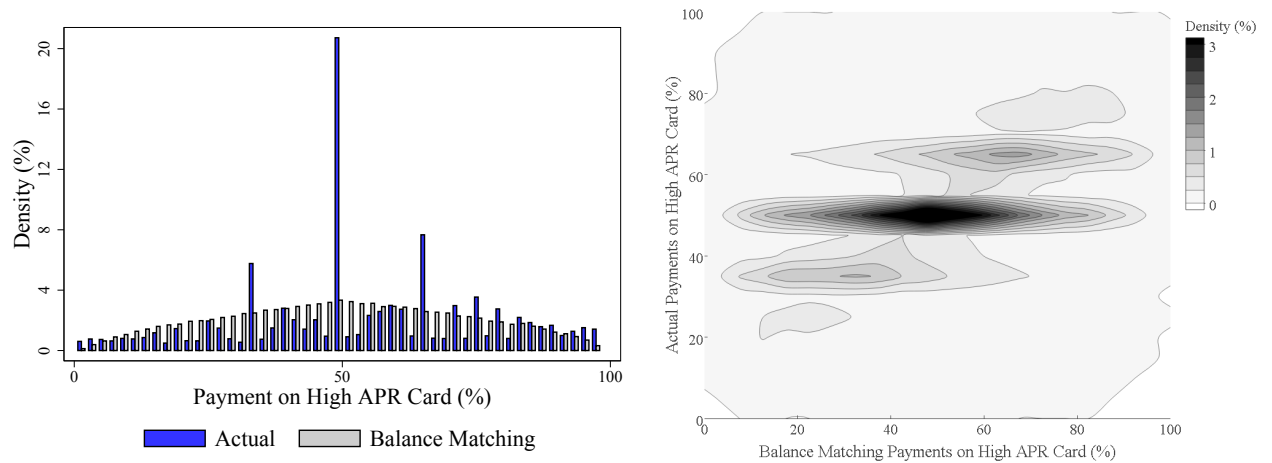
Figure A5: Example Credit Card Statement



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


Figure A6: Balance Matching

(A) Round Number Payment Sample



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


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



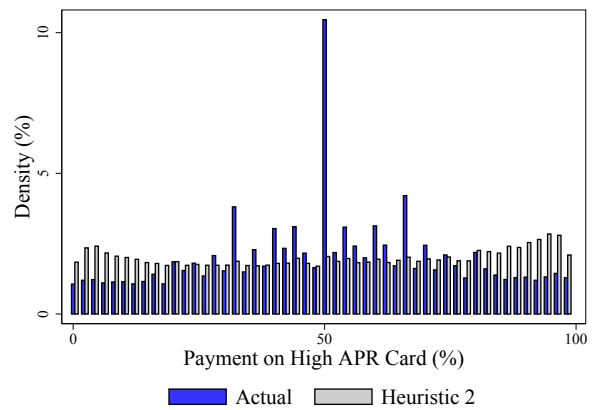
Note: Goodness-of-fit for different models of the percentage of payments on the high APR card. The left panel shows the Root Mean Square Error (RMSE), the middle panel shows the Mean Absolute Error (MAE), and right panel shows the Pearson Correlation Coefficient, which can also be interpreted as the square root of the R-squared. The round number sample restricts to observations where individuals make round number payments (multiples of £50), and the non-round number sample restricts to observations where individuals make non-round payment amounts (not multiples of £50). Random has repayments on the high APR card randomly drawn from a uniform distribution with support on the 0% 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 35 for details. The two-card sample is restricted to individual \times months in which individuals face an economically meaningful allocative decision. Goodness of fit is calculated using the 20% hold-out sample. See Section 2.2 for more details on the sample construction.

Figure A8: 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

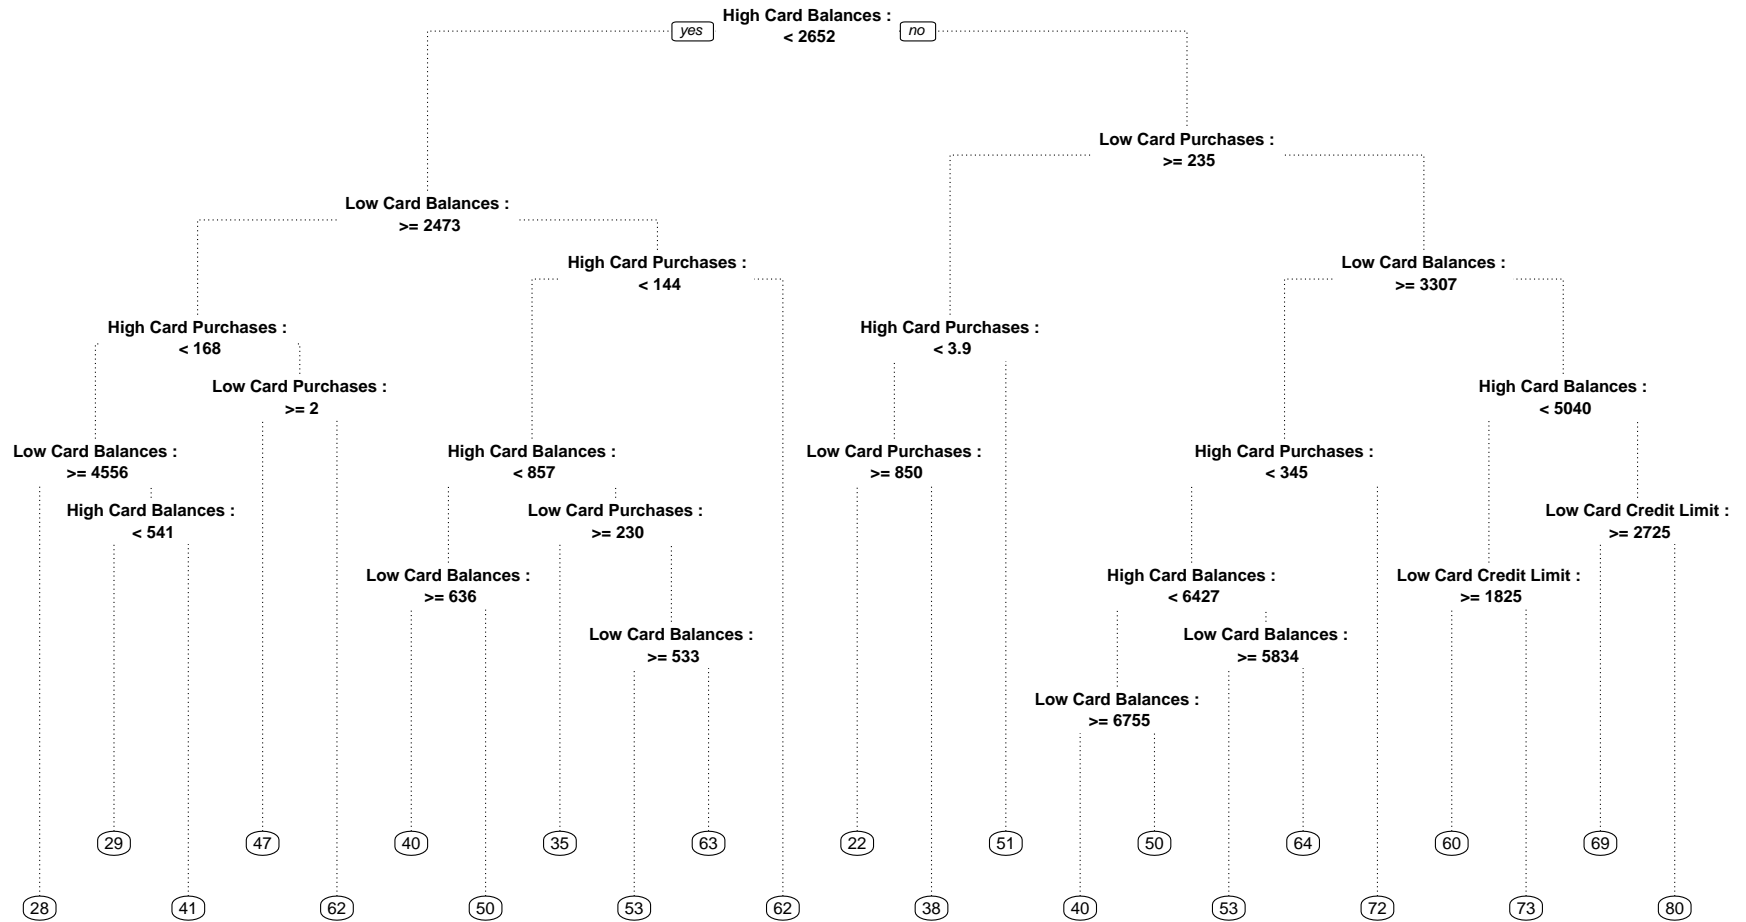


(D) Heuristic 4: Pay Down Lowest Balance



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 \times months in which individuals face an economically meaningful allocative decision. See Section 2.2 for more details on the sample construction.

Figure A9: High APR Card Payment Decision Tree




62

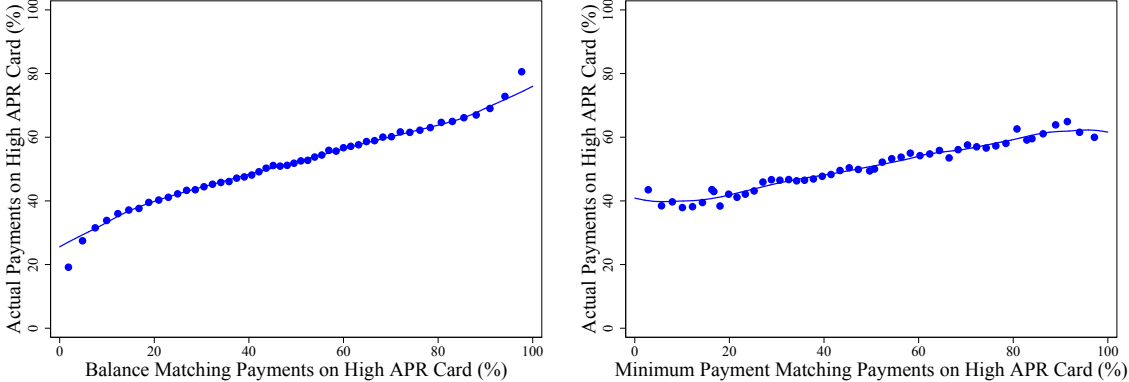
Note: Figure shows the decision (regression) tree for high APR card repayment. Top row is tree root. Nodes show the variable and split value at each branch. Bottom rows show predicted values at the end of each branch.

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

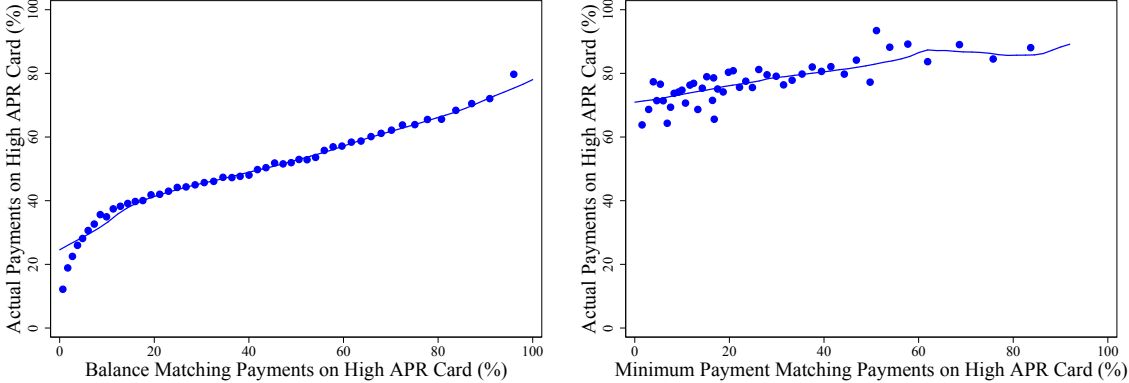
(A) Same Percentage Sample



(B) Different Percentage Sample




(C) Fixed Payment Sample



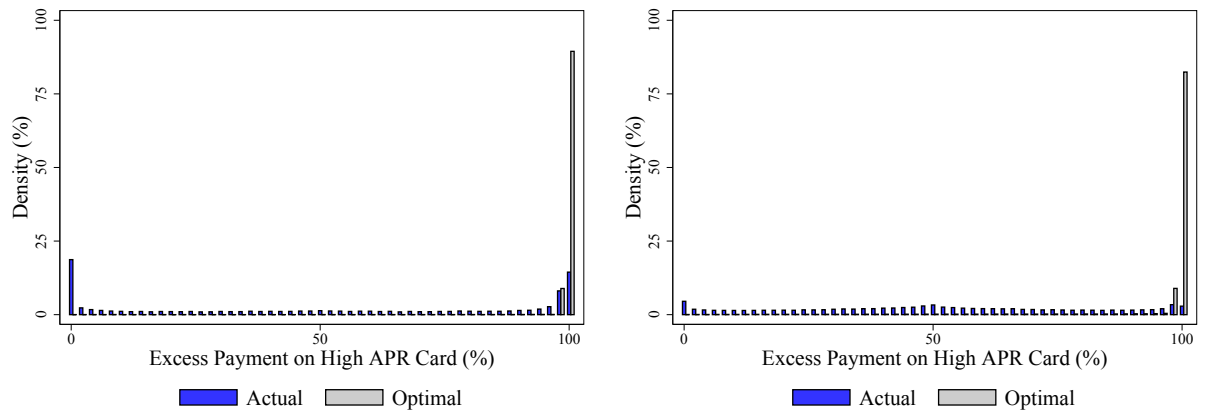
Note: Panels show binned-scatter plots of the actual percentage of monthly payment allocated to the high APR card (y-axis) and the percentage of total monthly payment allocated to the high APR card under the balance-matching heuristic (x-axis, left column) and minimum-payment-matching heuristics (x-axis,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

Figure A11: Actual, Optimal and Balance Matching Payments for Autopay (Left Column, 11% 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 \times months in which individuals face an economically meaningful allocative decision. See Section 2.2 for more details on the sample construction.

Table A1: Sample Restrictions

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

	Mean	Std. Dev.	10th	25th	Percentiles			90th
<i>i) As a % Total Monthly Payment</i>								
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	
<i>ii) Payments in £</i>								
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 \times months in which individuals face an economically meaningful allocative decision. See Section 2.2 for details.

Table A3: Annualized Interest Savings in the Unrestricted Sample

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

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

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

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

Note: Summary statistics for actual and optimal excess payments on the high APR card by quintiles of economic stakes and card age. Excess payments are calculated as the percentage of payments on a given card after subtracting out repayments needed to pay the minimum amounts due. Cells report mean values within the quintile. The two-card sample is restricted to individual \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 A5: Annualized Interest Savings Under Different Repayment Rules

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

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

Table A6: Goodness-of-Fit for Different Models

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

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

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

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

Note: Goodness-of-fit for different models of the percentage of payments on the high-APR card. The first column shows the Root Mean Square Error (RMSE), the second column shows the Mean Absolute Error (MAE), and third column shows the Pearson Correlation Coefficient, which can also be interpreted as the square root of the R-squared. Round and non-round samples are defined by whether repayments on the high APR card are multiples £50. See Footnote 35 for details. The two-card sample is restricted to individual \times 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 A8: Heterogeneous Types from 3-Way and 4-Way Horse Race Models

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

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

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

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

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

Table A10: Machine Learning Models Variable Importance

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

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

Table A11: Sensitivity Estimates Machine Learning Models Variable Importance

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

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

Table A12: Interest Savings from Optimal Dynamic Model

Panel (A): Optimizing Payments and Spending							
	Mean	Std. Dev.	Percentiles				
			10th	25th	50th	75th	90th
<i>Savings in £</i>							
3 Months	4.09	6.23	0.16	0.75	2.17	5.16	10.06
6 Months	15.25	20.03	0.66	2.97	8.64	20.53	38.20
9 Months	33.22	43.68	1.46	6.34	18.65	44.56	84.47
12 Months	57.78	76.53	2.55	10.80	32.01	77.39	148.74

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

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

Table A13: Correlations Between Payment Rules

Panel (A): Balance Matching vs. Min. Pay Matching			
	(1) Same Slopes	(2) Different Slopes	(3) Floor
Correlation	0.96 (0.00)	0.86 (0.00)	0.56 (0.02)
Panel (B): Balance Matching vs. Actual			
	(1) Same Slopes	(2) Different Slopes	(3) Floor
Correlation	0.63 (0.00)	0.41 (0.00)	0.50 (0.00)
Panel (C): Min. Pay Matching vs. Actual			
	(1) Same Slopes	(2) Different Slopes	(3) Floor
Correlation	0.61 (0.00)	0.28 (0.01)	0.23 (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 \times 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 \times 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 \times months in which the minimum payment determined by the floor value on both cards held by the individual, e.g. £25.

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

	(1) Both Cards Non-Autopay	(2) Both Cards Autopay
<i>i) Actual and Optimal Payments</i>		
Actual Payments (%)	51.21	51.11
Optimal Payments (%)	74.36	58.42
Actual - Optimal Payments (%)	23.15	7.30
<i>ii) Actual and Optimal Excess Payments</i>		
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 A15: Goodness-of-Fit for Different Models, Autopay and Non-Autopay Samples

	Both Cards Non-Autopay			Both Cards Autopay		
	(1) RMSE	(2) MAE	(3) Corr	(4) RMSE	(5) MAE	(6) Corr
<i>i) Main Models</i>						
Uniform Draw (0,100)	34.04 (0.19)	28.36 (0.19)	-0.01 (0.01)	36.90 (0.11)	30.30 (0.10)	-0.00 (0.00)
Optimal	36.40 (0.18)	30.65 (0.17)	0.25 (0.01)	32.86 (0.24)	20.81 (0.20)	0.35 (0.01)
1/N	17.64 (0.10)	12.63 (0.10)	-0.01 (0.01)	22.99 (0.06)	18.60 (0.07)	-0.00 (0.00)
Balance Matching	22.00 (0.14)	16.81 (0.11)	0.38 (0.01)	23.11 (0.13)	15.61 (0.09)	0.53 (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 \times 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 A16: Horse Races Between Alternative Models,
Autopay and Non-Autopay Samples

Panel (A): Uniform vs. Balance Matching		
	Both Cards Non-Autopay (1)	Both Cards Autopay (2)
<i>Win %</i>		
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)
<i>Win %</i>		
Balance Matching	75.21	61.02
Optimal	24.79	38.98

Note: Table shows percentage of individual \times 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 \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.