Even Self-Aware Consumers Are Overconfident

June 15, 2023

Matthew N. White[1] Christopher D. Carroll[2] Daniel J. Grodzicki[3]
David C. Low[4]

To explain why most consumers pay off "deferred-interest" credit programs early while many others fail to pay on time, we introduce a small (and empirically well-justified) tweak into an otherwise-rational model: consumers are only partially aware of their limited ability to execute future plans. In our structural model, most consumers are both "self-aware" (they realize they may accidentally fail the promotion) and "overconfident" (they underestimate their chances of doing so). So they choose payment plans that allow for some margin of error, but not as much margin as a perfectly self-aware agent would choose. Our model matches high early-exit rates, high failure rates, and many other important moments from the data.

html: https://llorracc.github.io/DefInt/
PDF: DefInt.pdf
GitHub: https://github.com/llorracc/DefInt

The views expressed in this paper are those of the authors and do not necessarily represent those of the Consumer Financial Protection Bureau, Johns Hopkins University, or the University of Delaware.

[1]White: Department of Economics, University of Delaware, and Econ-ARK
[2]Carroll: Department of Economics, Johns Hopkins University, ccarroll@jhu.edu, and NBER
[3]Grodzicki: unaffiliated
[4]Low: Office of Research, Consumer Financial Protection Bureau

1 Introduction

Economists have long wondered how agents who try to take account of their own limitations– who are "self-aware"– behave in a world too complex to fully understand.[1] Interesting evidence has been gathered in experimental settings, and in many environments boundedly rational rules-of-thumb are nearly as good as the omniscient solution. But it is hard to identify data, and develop models, to measure the prevalence and effectiveness of self-aware decision-making "in the wild."

In this paper, we present a model and evidence that together allow us to measure the extent of borrowers' self-aware behavior in the context of "deferred-interest" loans (specifically, credit card balances) that incur zero interest *if the borrower successfully pays off the entire balance by the end of the promotional period*, but incur very high interest if the balance has not been fully extinguished at the terminal date. For someone with no behavioral limitations, the optimal strategy is obvious: pay the minimum to avoid penalties until the last period, and extinguish the remaining balance at the end— the "bang-bang" solution. An optimizing borrower who was perfectly confident in their ability to execute a plan would choose the bang-bang plan.

Almost nobody does this. Instead, most people who "succeed" (i.e. pay zero interest) pay most or all of their balance well before the last period. In this paper we argue that even many of these *ex-post* successful participants paid their debt so quickly because they worried that they *might* make mistakes. This strategy accords with advice on navigating DI programs offered to us by ChatGPT:[2] "To give yourself some buffer, aim to pay off the loan a few months before the deferred interest period ends. This way, if any issues arise, you have extra time to resolve them without incurring interest charges."[3,4] A *perfectly* self-aware agent (i.e. one who knew their own exact error probabilities) would make a boundedly-optimal plan to do this.

This paper solves and estimates the first structural model of DI repayment behavior, and uses the model to understand the extent to which consumers in real-world data exhibit self-awareness. In line with ChatGPT's advice and our data, a large majority

[1] See, e.g., Herbert Simon (1950).

[2] Our exact query was: "I am worried that I might make a mistake paying off my deferred interest loan and have to pay penalty interest. How can I make a plan to avoid that?" Similar queries produce similar results.

[3] ChatGPT 4 accessed 2023-05-10 at 00:00; full exchange is in the appendix.

[4] Further ChatGPT conversation leads to the specific advice to "create an accelerated payment plan: instead of dividing the total balance by the number of months in the promotional period, divide it by a smaller number of months to calculate your accelerated monthly payment amount... to pay off the loan ahead of the deadline."

of agents in our model pay off their zero-interest debt earlier than necessary because they are aware of their own limitations. But agents are also overconfident: they do not pay their debt off as early as they would if they had *correct* beliefs about failure rates, and a significant number fail the program and end up paying high interest rates. This overconfidence essentially eliminates the benefits of the program to agents in our model on average. This provides rare and valuable evidence on the prevalence and implications of boundedly self-aware decision-making "in the wild."

In the model, agents begin with an exogenous amount of debt in a 12-month DI promotion; they make payments on this debt each month. In line with evidence we present, the model allows agents to choose from among a limited number of repayment plans, e.g. agents might choose to pay exactly $100 every month before month 12. Agents might make two kinds of mistakes when attempting to follow their repayment plans: (1) in any month they might forget to make a payment at all, and (2) they may fail to remember the distinct nature of month 12, and treat it like any other month. Consumers vary in both their actual and perceived probabilities of making these mistakes. To avoid the high costs of failing the promotion, self-awareness incentivizes consumers to create a buffer and pay their debt off early. But agents also value delaying repayment because their time preference rate is positive. Because DI payments generally constitute a limited fraction of a consumer's budget, we assume that borrowers are risk-neutral. To account for the fact that making one large payment at the end of the promotion may be difficult, we also allow for quadratic payment costs in the final period.

We discipline the model using unique data from a lender that offered a DI promotion. We have transaction-level information on each DI purchase and repayment, including the initial promotional purchase amount, monthly payments, account balances over time, and assessed fees and finance charges. We also have information on consumers' credit scores and self-reported income at origination. These data were matched to quarterly data from one of the three nationwide consumer reporting agencies ("NCRAs"), allowing us to observe cardholders' available credit, outstanding balances, delinquencies, etc on other credit products. Because our model focuses on consumers who are fairly sophisticated credit users with the ability to repay their debt, we drop consumers with revolving debt more than twice their income, as well as consumers with delinquencies within two years of the promotion. As intended, this results in an especially creditworthy sample: in our selected sample, the average credit score is 758 and DI promotional failure rate is 9.2%, while the numbers for the full sample are 737 and 14.4% respectively. We are thus likely to overestimate the extent of self-awareness, and underestimate the extent of overconfidence, among DI borrowers in general.

We estimate the model using the Simulated Method of Moments (SMM). The estimated model matches a wide variety of moments closely, including the size and persistence of payment amounts, missed payments, and promotion failure rates. Many moments used in estimation condition on borrower characteristics, including age, credit score, and income. This is important in part because it allows the model to account for the fact that borrower repayment behavior varies with (e.g.) credit score. In order to explain why so many DI borrowers pay off their debt so quickly, the estimated model needs most borrowers' perceived probabilities of making repayment mistakes to be substantial. But in order to match actual program failure rates, the estimated model also needs many borrowers' perceived mistake probabilities to be substantially lower than their actual probabilities. Hence, in the estimated model even self-aware consumers are overconfident. We also find a negative correlation between borrowers' perceived and actual mistake probabilities— consumers with higher probabilities of making mistakes (e.g. those with lower credit score) act as if their mistake probabilities are in fact particularly low.

Finally, we use the model to estimate the financial costs to DI borrowers of three kinds of deviations from standard "rational" consumer behavior: (1) positive mistake probabilities, (2) perceived mistake probabilities that differ from actual ones, and (3) cognitive idiosyncracies (e.g. a preference for rounding numbers). All three impose substantial costs even on the most creditworthy borrowers in our sample, but the costs are particularly large for less creditworthy borrowers.

2 The Deferred Interest Setting

Deferred interest (DI) credit card promotions are contracts that offer a zero annual percentage rate (APR) on a purchase or set of purchases that comprise a promotional balance for a fixed promotion period. Nevertheless, finance charges are calculated by the lender at a (typically higher than average) non-promotional APR from the time of purchase. If the promotional balance is not fully repaid at the end of the promotion period, the account is retroactively assessed this "deferred interest." Otherwise, no finance charged is assessed.

Zero interest promotions on credit cards became popular in the late 1990's and rose to prominence by the middle of the 2000's.[5] At this time, some of these promotions took the form of a DI contract, whereas others were "true" 0% APR promotions – without interest calculated during the promotional period.[6] In 2008, the Federal Reserve Board's (FRB) Office of Thrift Supervision (OTS) and the National Credit Union Administration (NCUA) jointly proposed new rules defining unfair, deceptive, or abusive practices (UDAP) for banks, thrifts, and credit unions. Made under auspices of the Truth in Lending Act (TILA) of 1968, the new rules applied to nearly all credit card issuers and declared deferred interest contracts to be unfair or deceptive.[7]

Within a few months of this decision, regulators reversed themselves and again permitted DI into the market. This action was taken under Regulation AA (2010) and required substantial enhancement to issuers' disclosures of APR and billing practices for consumers. In parallel, the Credit Card Accountability, Responsibility, and Disclosure (CARD) Act (2009) further codified apportionment of payments to balances on DI contracts (Sec. 104).[8] Interestingly, issuers did not lead the lobby to keep DI in the market. Rather, retailers argued that these products were vital to sustaining demand and that it was imprudent to stifle consumer spending in the midst of the Great Recession.[9] In light of new disclosure rules, and the championing of DI by retailers, issuers largely eliminated the product from their general purpose portfolios. Since 2010, DI has been offered almost exclusively on private label credit cards tied to specific retailers.[10] More recently, DI has become available to finance health care expenditures.[11]

Regulators argued for the removal of DI on two grounds. First, DI amounts to *"precisely the type of surprise increase in the cost of completed transactions that [the*

[5]See https://www.mymoneyblog.com/history-of-0-apr-interest-rates.html, which provides a graphic on the incidence of 0% APR usage using data from the Federal Reserve Board's Survey of Consumer Finances from 2001-2013.

[6]More specifically, these contracts provided 0% APR on revolving balances during the promotional period. Any balance remaining the end of the promotion would simply begin accruing interest at the "go-to" – non-promotional – APR.

[7]See https://www.govinfo.gov/content/pkg/FR-2008-05-19/pdf/E8-10247.pdf.

[8]This was also coupled with new disclosures. Ironically, the CARD Act also eliminated practices associated with contracts like DI, whereas DI is "carved out" of the CARD Act. Specifically, it eliminated "double cycle" billing (Sec. 102(a)) and retroactive increases in APR (Sec. 101(a)). This latter practice was commonly referred to as upward re-pricing (Consumer Financial Protection Bureau, 2015; Nelson, 2017).

[9]For details, see https://www.nclc.org/wp-content/uploads/2022/09/report-deferred-interest.pdf.

[10]Unlike general purpose cards, private label credit cards can only be used at specific retailers.

[11]See **https://www.consumerfinance.gov/about-us/blog/whats-the-deal-with-health-care-credit-cards-four-things-you-should-know** and **https://www.consumerfinance.gov/about-us/newsroom/cfpb-orders-ge-carecredit-to-refund-34-1-million-for-deceptive-health-care-credit-card-enrollment**.

new rule] is intended to prevent". This basis argues that DI harms consumers who are unaware of the deferred interest aspect of the contract.

Second, *"for the same reasons that consumers cannot ... reasonably avoid rate increases as a result of a violation of the account terms, consumers cannot ... reasonably avoid assessment of deferred interest as a result a ... failure to pay the balance in full prior to expiration of the deferred interest period ... disclosure may not provide an effective means for consumers to avoid the harm caused by these plans."* In other words, regulators argued that consumers are harmed by the intrinsic structure of the contract, whether or not they understand it. This latter justification alludes to other frictions endemic to consumer financial decisions:[12] They may not fully internalize the cost of using DI because they are not fully aware of their own ability to pay it off in time.

DI programs are an appealing setting in which to study consumer behavior for several reasons. First, they present a clean setting in which to model consumer repayment decisions. While the DI contract is technically open ended, the negative amortization implicit in its structure means that there is an optimal action for borrowers who believe they are perfect and inerrant.[13] Deviations from the minimum payment indicate the presence of repayment frictions, self-aware beliefs, or inherent behavioral biases.

Second, these contracts have drawn and continue to draw significant scrutiny from regulators as a form of unfair or deceptive practice on the part of credit card issuers. That in itself is a reason to study them, but until now there has been almost no academic attention to them. In addition, because they are seemingly designed to generate revenue from consumer mistakes, they can only exist in an environment in which such mistakes are prevalent.[14]

Third, this setting constitutes a fruitful empirical case study of the nature of consumer mistakes in a dynamic setting. In our data, we see consumers make repayments each

[12] See 74 Fed. Reg., 5498, 5528 (January 9, 2009) which can be found at https://www.govinfo.gov/content/pkg/FR-2009-01-29/pdf/FR-2009-01-29.pdf.

[13] That is, in the absence of any other financial constraints. Our empirical setting is designed to eliminate these as confounding factors to the extent possible. Absent other frictions, if the balance is not paid within the promotional period, the accrual of deferred interest can often mean the borrower owes more after the promotion than she originally borrowed, or that the balance has been negatively amortized. Note this is expressly ruled out for other card contracts in the CARD Act of 2009.

[14] However, it is important to note that one reason for these contracts being relegated to private label cards is that agreements between issuers and retailers endemic to these often include some back-end compensation to issuers from retailers. As aforementioned, retailers were the front line lobby for these products continuing existence. Additional profits DI brings to retailers are shared with issuers. There is very little publicly available information on these confidential agreements. Moreover, analysis of them is beyond the scope of this paper.

month. This unique aspect opens an avenue to inference of consumers' self-awareness and its evolution as they move through the promotion.

Misunderstanding of the DI contract has been highly cited as a cause of consumer harm both by regulators and the media. In line with this, significant effort has been put forward to enhance consumers' understanding of this product through a series of disclosure and apportionment rules.[15] Our analysis suggests that even consumers who understand the contract terms do not have the information (on real-world failure rates) necessary to make an informed choice about their repayment strategy.

3 Data

We now present a description of the full dataset and the several cuts required to ensure an appropriate sample, then motivate our structural model by graphically presenting the key stylized features of the data.

3.1 Data and Sample Selection

Our panel data comprise de-identified account-level information from a large lender's private label credit card portfolio. These data are part of the Consumer Financial Protection Bureau's (CFPB) Credit Card Database (CCDB). Account information cannot be tied to any particular consumer, nor can multiple accounts that belong to a single consumer or household be linked. As a result, our analysis is account-level rather than individual- or household-level. We focus on private label cards loans originated in a deferred interest promotion between January 2011 and September 2013. In what follows, we refer to this as our DI Sample.

We have transaction-level information on each DI purchase and repayment. This includes the initial promotional purchase amount, monthly payments, and all rates, fees and assessed finance charges. It also includes customers' reported income at origination.

[15] This is evidenced by the relegation of this product to private label cards. However, there has been very little empirical work on DI and consumer understanding. A notable set of studies is Consumer Financial Protection Bureau (2013, 2015). These analyses of the DI market have not found robust evidence that misunderstanding is the primary channel for consumer mistakes. One method relied on is repeat use by customers. The studies find that behavior of repeat-users is not substantially different from "first time" users, i.e. first time seen in the data. Notably, repeat use is a choice the consumer makes, whereby the population of repeat users is endogenously selected. Selection clearly affects measurement of learning by the "typical" consumer. Nevertheless, one might expect that, were learning primary, within consumer behavior would change with repeat use. The study finds little evidence of that.

Our model is of a DI borrower who opens and uses an account specifically to finance a single durable purchase using the DI promotion. We implement a series of sample restrictions to isolate an empirical sample of DI borrowers to align with this focus of the model. This allows the model to focus on cardholders' dynamic repayment behavior, rather than their purchase decisions.

We begin with the roughly 690k accounts for 12-month promotions. This is the modal promotion length in our data, and it corresponds to the promotion length in our model. We then drop 36k accounts with other promotions. Further dropping accounts with what appear to be data errors[16] leaves 620k accounts. We drop data from 19k accounts that pay off within the first two months; we view these borrowers as potentially outside our model, because opening a DI account and then closing it very quickly is not very different from not opening a DI account at all. To focus on accounts originated to purchase a product, we then restrict our attention to accounts originated within 15 days of the purchase date. This drops another roughly 76k accounts, leaving us with roughly 525k accounts. To obtain more variables we then match to these data to the broader CCDB; we drop the roughly 16k accounts that fail to match.

DI promotions often apply to only a single purchase. Therefore, we then drop accounts used to purchase more than one item; this is a major restriction and drops roughly 289k accounts, leaving us with roughly 220k accounts. Then we drop further problematic obervations through a further series of data cleaning steps. The most significant of these is that we drop roughly 12k accounts that ended the promotion with a positive balance and yet were labeled as succeeding the promotion; we suspect these accounts were not actually in the 12-month promotion. Because of these and several other data cleaning steps we are left with 203k accounts.

These data were matched to quarterly data from one of the three nationwide consumer reporting agencies ("NCRAs"). These data allow us to observe cardholders' outstanding balances, available credit, delinquencies, etc. on other tradelines, along with their age and credit score. About 2k accounts cannot be matched to the NCRA data; another 19k accounts are missing important variables in the NCRA data that we use later, so we also drop them.

Our model focuses on cardholders that are sophisticated credit users with the ability to repay their debt. To reflect this focus, we drop accounts that were 30 days delinquent

[16]The vast majority of these problematic observations come from accounts with well over twelve months of data and yet no indication that they failed the promotion. We suspect these accounts were misclassfied and not actually in the 12-month promotion.

on debt within two years of the start of the promotion. Of the remaining 181k accounts, this drops about 51k. We further restrict the sample to accountholders between the ages of 25 and 65 at the time of purchase, leaving us with roughly 107k accounts. Finally, we drop accounts with revolving debt exceeding twice their monthly income; this leaves us with 84k accounts.[17]

We implement these sample selection steps for internal, not external, validity. However, this sample is highly informative of the overall market, given that prime consumers account for 90% of DI purchase volume.[18] Table 1 below compares summary statistics for the full and selected sample. As intended, our restrictions leave us with a more creditworthy sample.

Table 1 Sample Summary Statistics

Variable	Full Sample	Selected Sample
Credit Score	737	758
Income ($)	80,319	91,096
Credit Limit ($)	2872	3364
Purchase Amount ($)	1487	1533
Accumulated Deferred Interest ($)	175	169
Months in Promotion	9.6	8.4
Failed Promotion (%)	14.4	9.2
Observations	688,567	84,397

NOTE: Table provides average statistics for the full and selected samples.

3.2 Stylized Facts to Inform Our Model

Figure 1 plots the distribution of the number of months consumers spend in the promotion. Two facts stand out. First, it is very common for consumers to exit the promotion early: about two-thirds of consumers exit before month 11 and more than half exit before month 10. Thus, even among borrowers that chose to take up a DI product, most decide before the end of the promotion that staying in it is not worthwhile. Second, among those who do stay in the promotion to month 12, almost *half* fail the promotion.

[17]Credit reports generate point-in-time information on credit balances. This makes it difficult to disentangle true revolving debt from revolving balances not incurring finance charges, e.g. in their grace period. As a result, it is difficult to differentiate heavy credit card *users* from heavy credit card *borrowers*. However, excessively high balances, such as those we eliminate here, are more likely to arise from true debt.

[18]https://files.consumerfinance.gov/f/201512_cfpb_report-the-consumer-credit-card-market.pdf

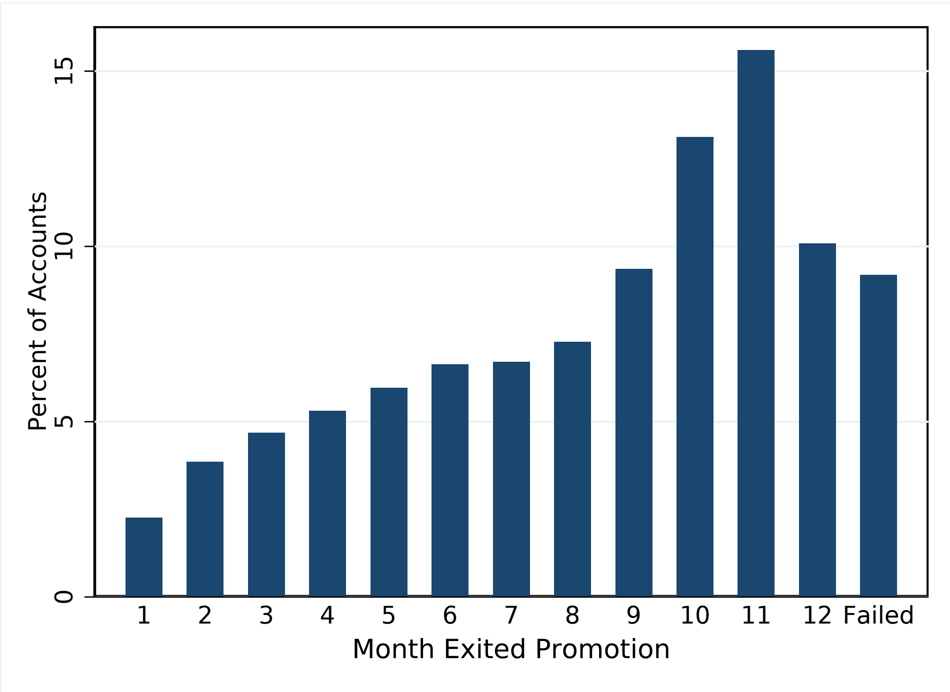

Notes: The figure shows the percent of accounts that exit the promotion in each month.

Figure 1 Distribution of Months Spent in Promotion

While staying in the promotion as long as possible would appear to be the best way to maximize the returns from it, this strategy in fact appears to be very risky. This risk in turn may justify the decision of so many consumers to leave the promotion early.

A second important fact is that borrower payments are persistent over time. Figure 2 plots the distribution of the change in payment from the second month to the penultimate month, by the number of months an account spends in the promotion. Regardless of the month an account pays off, the median borrower makes the exact same payment in the penultimate month as they did in the second month. Even at the 25th and 75th percentiles, the penultimate payment and the second payment are quite similar, indicating that many consumers have payment plans to which they closely adhere. This is especially evident among accounts that pay off between months 9 and 12, for which the distributions in payment changes since month 2 are especially compressed. Among borrowers that successfully pay off their balance between months 9 and 12, most appear to have done so by making consistent and roughly equal monthly payments. Perhaps most surprisingly, even among borrowers that fail the promotion, half of borrowers make the same or larger payment in month 12 as in month 2.[19] This suggests that most

[19] Among borrowers who fail the promotion, the 25th percentile of payments in month 2 is $35 and

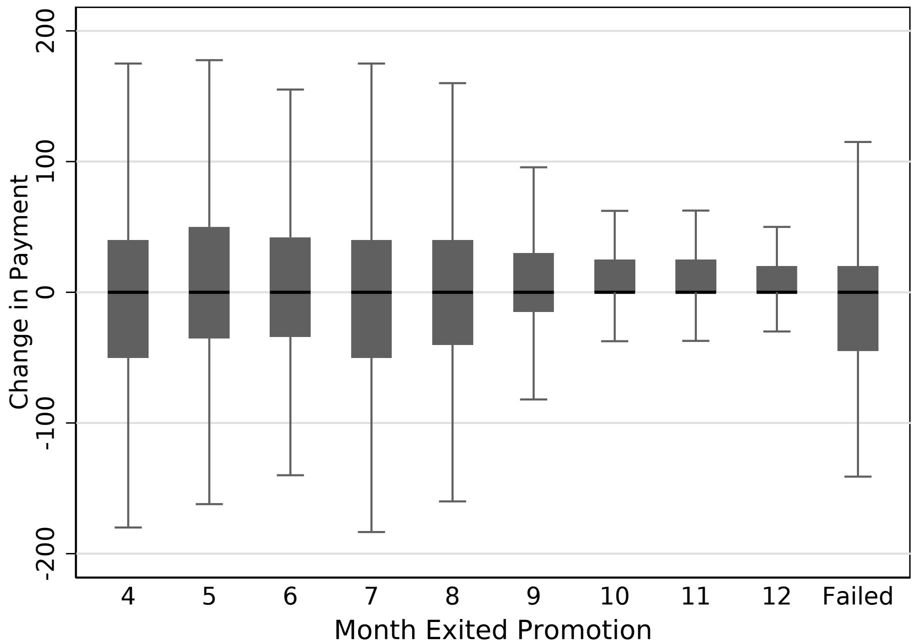

Notes: The figure shows box-and-whisker plots of the distribution of the dollar change in payments between the second and penultimate month, by the month an account leaves the promotion. For those who fail the promotion, the "penultimate" month is defined as month 12. Dark lines are medians. Gray boxes extend from 25th to 7th percentiles. Whiskers extend from 10th to 90th percentiles.

Figure 2 Payment Changes Between Second and Penultimate Month, by End Month

borrowers who fail the promotion do so on "autopilot", i.e. not by forgetting to make a payment nor by intentionally making a particularly small payment, but rather by sticking to their normal payment plan and failing to adjust for the unique nature of month 12.

Figure 3 examines the payment choices that lead to these patterns by categorizing borrowers' non-terminal payments. Minimum payments are rare; they comprise roughly 10% of payments early in the promotion, but even in the month that borrowers' fail the promotion (by making a non-terminal payment in month 12) they represent less than 30% of payments. However, persistent payment heuristics are common. In every month about 60% of payments are close to the borrowers' modal payment; of these around half are within \$5 of the original debt divided into either 9, 10, 11, or 12 equal monthly payments, while most of the rest are divisible by \$50. This suggests that borrowers choose among a fairly small number of repayment heuristics and change heuristics rarely.

---

the median payment in month 2 is \$75. Thus most of these borrowers pay more than the required minimum payment every month.

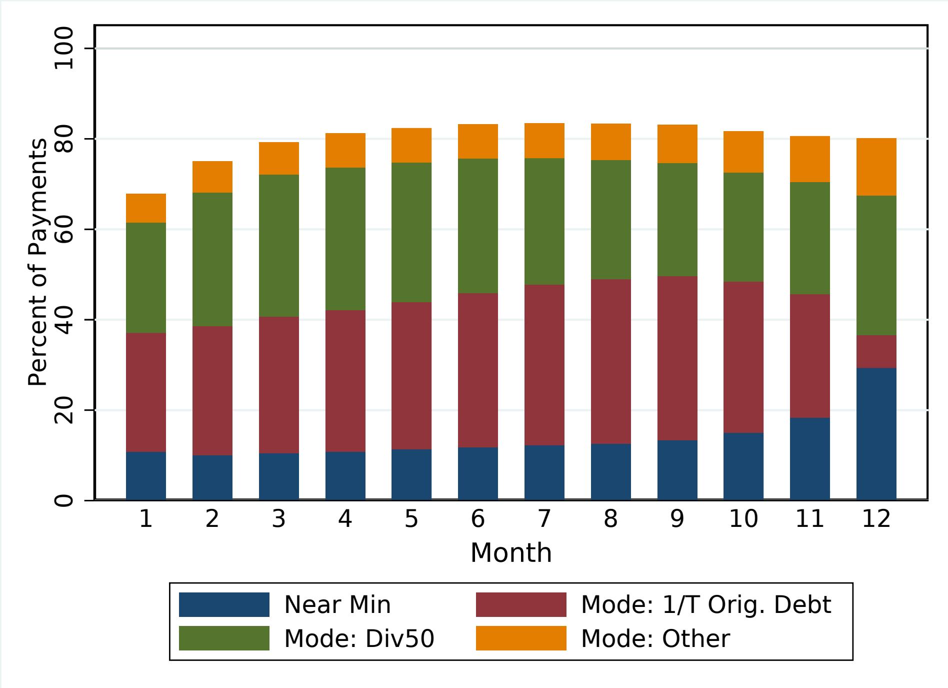

Notes: The figure shows the percent of non-zero, non-terminal payments in each month that fall into various categories. "Near minimum" payments are within \$10 of the minimum required payment. "Modal" payments are defined as within \$50 of a borrower's modal payment. "Mode: 1/T" are modal payments that are also within \$5 of the original balance divided by 9, 10, 11, or 12. "Mode: Div50" payments are modal payments that are exactly divisible by \$50. "Mode: Other" are other modal payments. When payments could fall into multiple categories, they are placed into the earliest one discussed.

Figure 3 Categories of Non-Terminal Payments

To investigate why borrowers fail the promotion in more detail, Figure 4a plots the distribution of unused credit on other cards for those who fail the promotion; Figure 4b does the same, but normalizes unused credit on other cards by consumers' promotional balance in month 12. In particular, the latter shows that over two thirds of borrowers who fail the promotion have at least five times their promotional month 12 balance unused on other credit cards. This provides more evidence that consumers in our sample do not fail the promotion simply because they are liquidity constrained.

Finally, Figure 5 provides information on purchase amounts in our sample. Figure 5a plots the distribution of purchase amount divided by borrower income. DI borrowers generally spend 2-3% of their annual income on their DI purchase, an amount large enough to be meaningful but likely too small to place significant financial hardship on most of the creditworthy consumers in our sample. Figure 5b plots the distribution of purchase amount divided by available credit on other credit cards. For over three-quarters of our sample, the DI purchase represented less than 20% of unused credit

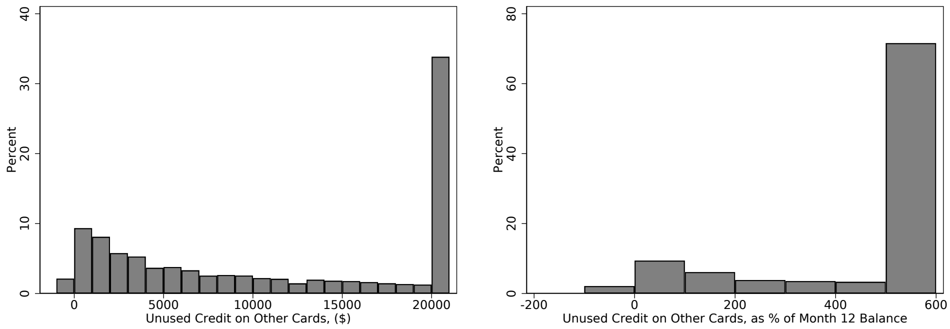

(a) Dollars **(b)** Percent of month 12 balance

Notes: The figure on the left shows the distribution of unused credit available at the end of the promotion on other credit cards among accounts that failed the promotion (winsorized above at \$20,000). The figure on the right shows this same information expressed as a percent of the remaining promotional balance in month 12 (winsorized above at 500%).

Figure 4 Unused Balances on Other Cards, for Borrowers Who Failed the Promotion

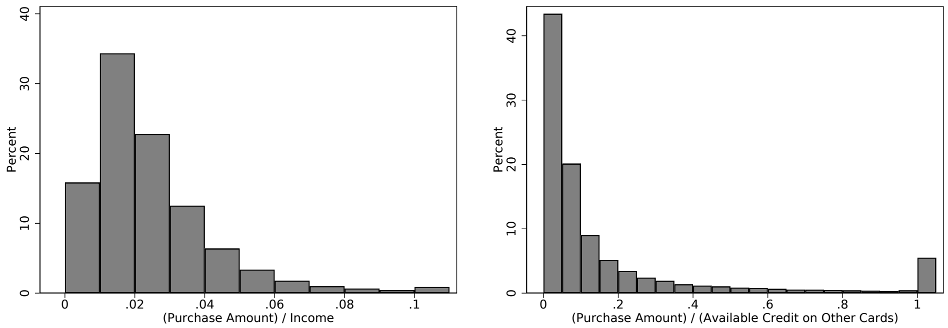

(a) (Purch. Amt)/ Income **(b)** (Purch. Amt)/(Other Credit)

Figure 5 Distribution of Normalized Purchase Amount

Notes: The figure on the left shows the distribution of the purchase amount expressed as a fraction of borrower income (winsorized above at .1). The figure on the right shows the distribution of the purchase amount expressed as a fraction of available credit on other credit cards at the time of purchase, winsorized above at 1.

available on other credit cards. This suggests that for most borrowers in our sample the DI cards themselves were not necessary to purchase the items, so for these borrowers the benefit of the DI product was the promotional interest rate.

4 Model

This section presents our structural model of the behavior of individuals who hold deferred interest debt. During a fixed promotion period, the individual does not owe any interest payments on this debt. However, the interest that *would* ordinarily be owed on credit card debt accrues (and compounds) in a separate account; if the individual does *not* pay off the entire balance of the debt before the end of the promotion (one year), he must pay the accumulated deferred interest (along with the remaining balance).

This is the first model of deferred-interest repayment behavior that we are aware of, so we have no direct precedent to follow. To construct our model we begin with a fairly standard framework of a rational, forward-looking consumer managing debt payments. As detailed below, we must make several adjustments to this framework to allow the model to match the data, but these adjustments (e.g. rounding) are already well-accepted in the household finance literature. We use this standard framework to demonstrate that, even if consumers are approximately rational, they still must substantially underestimate their chances of failing the deferred-interest promotion to make the payment choices they do.

4.1 Account Mechanics

Let discrete time be indexed by t, labeling $t = 0$ as the start of the promotion; the promotion lasts T periods, with any remaining balance due to be paid at the start of period T. Denote agent i's outstanding deferred interest debt as D_{it}, and their accumulated deferred interest as Z_{it}. Each period, the agent makes a payment of P_{it} on the deferred interest debt, paying down the debt: $D_{it+1} = \max(D_{it} - P_{it}, 0)$; if there would be $D_{it+1} \leq 0$, the debt has been successfully paid off and any overpayment is instantly returned to the agent. If $D_{iT} > 0$, then the individual is deemed to have "failed the promotion" and must pay Z_{iT} to the credit card issuer, along with the remaining debt D_{it}. During the promotion, deferred interest accrues at the per period interest rate of r_Z, so that $Z_{it+1} = Z_{it} + r_Z(Z_{it} + D_{it})$. If a payment is *missed* $P_{it} \leq \underline{P}$, not paying at least some minimum level, then a missed payment fee of $M > 0$ is immediately charged to the borrower.

Motivated by evidence presented in section 3, the individual does not get to choose P_{it} in every period; rather, he chooses a *payment heuristic* n_{it} from a finite menu N. There are two types of heuristics: *timing-based* and *dollar-based*. Choosing a timing-based heuristic $0 < n_{it} <= T$ at time t means that the individual *plans* to pay $\widetilde{P}_{i\tau} \approxeq D_{it}/n_{it}$ in this and subsequent periods τ. Choosing the special heuristic $n_{it} = 0$ means the agent plans to make the minimum allowed payment on the debt in all periods, $\widetilde{P}_{i\tau} = \underline{P}$. Alternatively, choosing a dollar-based heuristic $n_{it} \gg T$ at time t means that the agent intends to make payments of exactly $\widetilde{P}_{i\tau} = n_{it}$ until the debt is paid off. [20] Every repayment heuristic includes an intention to pay all remaining debt in the final period $t = T - 1$; model borrowers never intend to fail to the promotion.

The agent's heuristic is only a plan, and he will not necessarily execute this plan to perfection; rather, his payment behavior is subject to two kinds of errors that we observe in the data. First, he might miss any given payment with probability ς_i (representing his likelihood of a *screw-up*), resulting in a realized payment of $P_{it} = 0$ in that month. Second, if he reaches period $T - 1$ with remaining debt $D_{iT-1} > 0$, the agent might *fail to notice* it is the final period of the promotion and make his ordinary planned payment $\widetilde{P}_{it}$ rather than paying all remaining debt. Conditional on the borrower's observable characteristics, this more serious mistake occurs with probability φ_i, with an analogously escalated mnemonic. Not noticing it is the final period of the promotion is a (potentially) very costly mistake: if his ordinary payment does not cover the remaining debt, then the borrower will fail the promotion and he will have to pay the accumulated deferred interest Z_{iT}.

Whenever he chooses a heuristic, the agent believes he is making a *once and for all* decision and will never get to make a different choice. However, he will actually have the opportunity to choose a new heuristic with probability π in each succeeding period. When such an opportunity arises, the borrower might choose to remain on the same repayment heuristic, switch to another one, or simply pay off all remaining debt (by choosing $n_{it} = 1$ to pay off in a single month).

4.2 Observed and Unobserved Heterogeneity in Mistake Probabilities

As noted above, an agent's probability of missing a payment ς_i and of failing to notice month 12 φ_i are individual-specific. In this subsection we discuss how we discipline this heterogeneity, which is unobservable in our data, by using heterogeneity in observables.

We observe several demographic characteristics of borrowers at account origination. We collect several of these – a borrower's age j_i, FICO score s_i, (log) monthly income y_i, and percentage of available credit used c_i – into an observed vector x_i for each borrower.

We assume that each agent's true probability of missing a monthly payment ς and failing to notice the final month φ_i are determined by the normal distribution's sur-

[20]The values for dollar-based heuristics all far exceed the number of months in the promotion; using n_{it} to represent both "months to pay off" and "dollar amount to pay" is merely for convenience.

vivor function on a linear function of observable data, with an additional unobservable component for missed payment probability. The coefficients for the vector of observable characteristics are given by α and μ.

$$\varsigma_i = 1 - \Phi(\alpha \cdot x_i + \epsilon_i^s), \qquad \varphi_i = 1 - \Phi(\mu \cdot x_i). \tag{1}$$

Failing to notice that the final month of the promotion is (at most) a one-time event for each borrower, and this event is determined by the draw of the random variable ϵ_i^f:

$$\mu \cdot x_i + \epsilon_i^f < 0 \Longrightarrow \text{fail to notice it's the final month.} \tag{2}$$

The random terms ϵ_i^s and ϵ_i^f are normally distributed with zero mean. ϵ_i^f has unit variance. The standard deviation of ϵ_i^s is σ_α, which represents the extent to which unobserved characteristics affect missed payment probability. These two error terms are drawn from a bivariate normal distribution with correlation coefficient $\rho_A \in (-1, 1)$.

A primary goal of our model is to estimate if agents' perceived probabilities of missing payments and failing the promotion align with their actual probabilities of doing so. Therefore, we assume that borrowers have possibly inaccurate *beliefs* about these probabilities. We further assume that a borrower's perceived probability of failing to notice month 12 is binary; i.e. a consumer is either certain he will notice it is month 12, or certain he won't. We make this assumption because non-binary beliefs cannot be identified in our model; the costs of failing the promotion are so high, relative to any plausible benefits of continuing past month 12, that agents choose payment heuristics that do not require them to notice it is month 12 if they perceive any substantial probability of failing to notice month 12. We view this binary perceived probability as a shorthand to differentiate between borrowers with even vaguely correct beliefs about critical failure versus those who *act as if* it is essentially impossible. We denote agent i's perceived probability of missing payments and failing the promotion at time t as, respectively, $\widetilde{\varsigma}_{it} \in [0, 1]$ and $\widetilde{\varphi}_{it} \in \{0, 1\}$.

$$\widetilde{\varsigma}_{it} = 1 - \Phi(\beta \cdot x_i + \xi_i^s), \qquad \widetilde{\varphi}_{it} = \mathbf{1}((\nu + t\delta) \cdot x_i + \xi_i^f < 0). \tag{3}$$

Note that perceptions of the probability of missing a payment are generated from the same survivor function transformation as the actual probabilities, but with different coefficients and error terms. These coefficients and error terms are static throughout the promotion. However, δ induces a drift in agents' perceived probabilities of failing to notice month 12; this represents the possibility that the borrower switches from being fully optimistic about the prospect of noticing the final month to fully pessimistic (or vice versa). Like the error terms ϵ_i^s and ϵ_i^f, we assume that each borrower's "perceived error terms" (ξ_i^s and ξ_i^f) are drawn from a bivariate normal distribution with mean zero, standard deviations of σ_β and one (respectively), and a correlation coefficient of ρ_B.

4.3 Borrowers' Preferences

Model borrowers are risk-neutral expected utility maximizers. The reason they hold deferred-interest debt in the model is because they prefer to repay their debt in the future rather than immediately. To parameterize this benefit, we assume that, relative to paying off all of the debt right away at $t = 0$, an agent who holds D_{it} dollars of debt from period $t-1$ to t gets a utility flow of rD_{it}. This "time-preference interest rate" represents an unspecified combination of inflation (the debt is in nominal dollars), financial returns to saving, pure time preference, and risk aversion or the benefits of flexible payment timing.

Borrowers in our model experience frictions associated with participating in the promotion, which we capture as adjustments to the value of each heuristic. Dollar-based heuristics provide utility depending on how round that dollar value is (i.e. in increments of \$100 or \$50.) Moreover, participating in the deferred interest promotion is burdensome, relative to simply paying the debt all at once: the borrower must think about his repayment plan, and then take the time to make each payment (as auto-pay is not available). We capture this cost of more active participation with the parameter τ_1, which formally is a bonus to utility from choosing $n_{it} = 1$ to pay all of the debt this month; it should be interpreted as the effort cost of *any other plans*. The bonus utility for a heuristic is thus:

$$\kappa(n_{it}) = \begin{cases} \tau_1 & \text{if } (n_{it} = 1) \\ \kappa_{100} & \text{if } (n_{it} \gg T) \ \& \ (n_{it} \equiv 0 \pmod{100}) \\ \kappa_{50} & \text{if } (n_{it} \gg T) \ \& \ (n_{it} \equiv 0 \pmod{50}) \\ 0 & \text{otherwise.} \end{cases} \tag{4}$$

If more than one condition is fulfilled, the higher one on the list is used.

Although we model borrowers as risk-neutral, this is because we only study a small segment of agents' overall financial behavior. In a model with a broader scope, agents would be risk-averse; in particular, making a large payment in the final month might be costly. To capture this aspect of reality, we include a penalty term in borrower utility. Specifically, if the realized payment in the final month P_{iT-1} exceeds the borrower's usual planned payment $\widetilde{P}_{iT-1}$, the borrower experiences a quadratic utility cost, scaled by coefficient ω:

$$\omega \max(P_{iT-1} - \widetilde{P}_{iT-1},\ 0)^2 \tag{5}$$

At the beginning of the promotion, each borrower draws and observes a preference shock η_{in} for each payment heuristic, idiosyncratic to themselves and unobserved to the econometrician; they are normally distributed with mean zero and standard deviation σ_η. Each borrower's vector of preference shocks represents heterogeneity of preferences from sources that are fully outside of the model's scope. These preference shocks are not driven by beliefs or preferences, but instead represent a structural error term. Because agents are risk neutral, both these shocks and the shifters above can be considered to be in dollar-equivalent utility units.

At any time t when the agent can choose a repayment heuristic, he computes the present expected value of each plan $n \in N$ and chooses the plan with the highest expected value. This computation includes the expected value of delaying repayment, fees paid for missed payments, the deferred interest penalty, and the quadratic utility cost of making a large final payment, as well as the relevant preference shifter and idiosyncratic shocks.

To express this expected value formally, first denote the individual i's realized missed payment fee in period t as:

$$\widehat{M}_{it} = \mathbf{1}((P_{it} < \underline{P})\ \&\ (D_{it} > 0))M. \tag{6}$$

That is, the borrower pays fee M only if his payment is below the minimum level and he has remaining debt. The agent's subjective value from choosing a heuristic can then be expressed as:

$$V_t(n_{it}; D_{it}, Z_{it} \mid \widetilde{\varsigma}_i, \widetilde{\varphi}_{it}) = \tag{7}$$

$$\mathbb{E}_t \left[\sum_{\tau=t+1}^{T} \left[rD_{i\tau} - \widehat{M}_{it} \right] - \omega \max(P_{iT-1} - \widetilde{P}_{iT-1},\ 0)^2 - \mathbf{1}(D_{iT} > 0) Z_{iT} \middle| \widetilde{\varsigma}_{it}, \widetilde{\varphi}_{it} \right] + \kappa(n_{it}) + \eta_{in}.$$

Note that the subjective expectation holds fixed the borrower's perception of his error probabilities. Borrowers do not anticipate changing their beliefs or their payment plan in the future.

5 Estimation

We estimate the parameters of the structural model using the simulated method of moments (SMM), seeking to minimize the weighted difference between simulated model outcomes and their empirical counterparts. This section first describes calibrated parameters and summarizes the objective function procedure, then specifies ways in which the simulated and empirical data are "sliced" to condition for individual moments. We then present and discuss the set of moments used in the estimation, and how they identify the structural parameters.

5.1 Calibrated Features

Several features of the model are set exogenously, rather than estimated. Most basically, we set the length of the model to $T = 12$ periods, for the annual promotion used in the dataset, and the deferred interest rate has an APR in excess of 20%.

We also incorporate two minor behavioral quirks to match secondary features of the data. First, a "rounding threshold" is randomly drawn whenever a timing-based heuristic is selected, so that actual payments might be rounded to the nearest penny, \$1, or \$5. This has no substantive effect on the pace of repayment. Rather, it allows the simulation to match the "micro-rounding" feature observed in the data, and generates a small number of accounts who pay *just short* of one-twelfth of the debt each period, and then fail the promotion with a very small remaining debt, as in the data. Second, whenever a model borrower misses one or more payments, there is a 50% chance that they will include all missed payments in their next non-zero payment. In the data, about half of missed payments are made up in this way, appearing to be *late* (and thus included in the next account month) rather than actually missed. Such late payments are usually (but not always) charged the missed payment fee; we assume that agents believe the fee will be charged for missed or late payments.

In the dataset, most borrowers cluster around a small number of repayment strategies. First, about 45% of all non-terminal payments are in even increments of \$100; another 15% are in increments of \$50, while much smaller percentages are rounded to \$10 or \$25. We thus include in the menu of heuristics all increments of \$50 up to \$400.[21] Second, a noticeable proportion of accounts pay *almost exactly* one-twelfth of the original debt in any given period– about 16% of payments. A smaller share of payments are very close to one-tenth of the original debt, but no other fractions have more than a trivial share

[21]To-do: Allow \$5 increments and increase cap to \$600.

of payments. We interpret this as borrowers considering the possibility of making even payments over the twelve month promotion, or dividing by ten as an easy mathematical strategy, and hence include $n_{it} = 10$ and $n_{it} = 12$ as available heuristics. To allow model borrowers to suddenly exit the promotion by paying all remaining debt, we also include $n_{it} = 1$ in the menu. Finally, the minimum payment plan n_{it} is available for borrowers to pay exactly $\underline{P}$, which is \$20 in the promotion for our data.

Model borrowers derive a benefit from the deferred interest promotion by delaying repayment, at the risk of potential missed payment fees and deferred interest. Under the assumption that they are preference maximizers, conditional on their subjective probability beliefs, the time preference rate of return is a key determinant of model behavior. In the estimation, we identify the distribution of borrower beliefs by (effectively) inverting the relationship: seeking belief parameters that justify observed behavior. If the time preference rate r on the order of financial returns available to the borrower in the short run (say, 1% or lower, annually), the potential benefits of the promotion would be *very* small. To justify borrowers' participation in the promotion, model agents would need to believe that they could *almost never* make any kind of mistake the promotion, as they're risking 20-40% of the original debt in deferred interest in exchange for less than 1% of the debt in "delay value"– a poor proposition. We trim the empirical sample to exclude borrowers with any recent adverse credit events, or who potentially don't have access to other credit options, implying that liquidity constraints are unlikely to bind for most or all borrowers. Nevertheless, we calibrate the time preference rate at 5% annually ($r = 1.05^{1/12} - 1$) to give model agents the benefit of the doubt about their potential gains from the deferred interest offer. As will be revealed by the estimated model, even with such a high value of delaying repayment, model borrowers must hold *unrealistically overconfident* beliefs about their propensity of making mistakes in order to rationalize observed repayment behavior. With a lower value of r, the model would require agents to have even more implausible beliefs.

5.2 Objective Function

The goal of the estimation is to find the vector of structural parameters θ^* under which the repayment behavior of simulated model agents most closely resembles that of the real borrowers in the deferred interest data. As normal in structural estimation, we specify a set of conditional statistics (moments) of the data, and seek to match model-generated moments to their empirical counterparts.

Let there be M real-valued moments to match and K structural parameters to estimate, with $M \gg K$, and denote the vector of empirical moments as $\Gamma \in \mathbb{R}^M$. For a given guess of the structural parameter vector $\theta \in \mathbb{R}^K$, computation of the objective function proceeds in four steps:

1. Find the model borrower's (subjective) value of each repayment heuristic $n_{it} \in N$ as a function of the state variables $\widetilde{\varsigma}_{it}$, $\widetilde{\varphi}_{it}$, D_{it}, Z_{it} in each period of the promotion t, ignoring idiosyncratic preference shocks η_i.

2. Simulate several copies of each borrower in the data, starting from their observed initial conditions (characteristics and original debt) and drawing idiosyncratic preference shocks over heuristics. During the simulation, model borrowers choose their optimal repayment plan and try to execute on it for T periods, potentially missing payments and being surprised by the chance to choose a new plan.

3. Calculate the set of moments on the simulated repayment data, following the same procedure as for the empirical moments.

4. Compute the difference between the vectors of simulated and empirical moments, then take its inner product through an $M \times M$ weighting matrix Ω.

Denote the functional mapping from parameter vector θ to the set of simulated moments as $g(\theta)$, as described in steps 1-3. The objective of the estimation is to minimize the weighted difference between simulated and empirical moments (step 4):

$$\theta^* = \arg\min_{\theta \in \mathbb{R}^K} \, (g(\theta) - \Gamma)'\Omega(g(\theta) - \Gamma). \tag{8}$$

The weighting matrix Ω is constructed as the inverse of the bootstrap covariance matrix for the empirical moment vector Γ. That is, we re-sample the dataset 1000 times, drawing(with replacement) borrower observations from the original set and computing empirical moments on each pass, then calculating the column-wise covariance of the $1000 \times M$ matrix of bootstrapped moments. This weighting matrix controls for the scale of the moments (so that large numbers are not unduly given more value in the estimation) while putting more weight on moments that the data is "confident" in– those with low variance. Moreover, statistically related moments are appropriately downweighted, as their values do not provide fully independent information about the data.

The minimization procedure was conducted using a combination of quasi-Newton search and a standard Nelder-Mead polytope method. When employing Newtonian methods, the Hessian of the objective function was approximated as two times the inner product of the Jacobian of $g(\theta)$ through the weighting matrix W; this effectively computes the Hessian as if $g(\theta)$ is locally linear. In preliminary estimations, blocks of parameters were estimated on subsets of the moments (setting rows and columns of Ω for other moments to zero), motivated by the identification arguments presented below. Once most of the simulated moments were reasonably close their empirical counterparts, the minimization procedure was performed on all parameters simultaneously, with the full weighting matrix. Standard errors for the estimated parameters were calculated based on the inverse of the approximated Hessian.

5.3 Categorization of Borrowers and Payments

Many of the moments used in the estimation are conditioned on a particular range of a borrower characteristic, partitioning the domain of each variable; we refer to these as "observable groups". As our estimation dataset excludes both young and old borrowers, we divide accounts into four age groups of roughly equal size: 32 and under, 33 to 42, 43 to 52, and 53 and over. Due to our restriction to borrowers without recent derogatory credit report entries nor delinquent debt, the distribution of credit scores in the estimation data is significantly higher than the general population, and tightly clustered. We use five credit score groups that approximately correspond to the quintiles of the data: 705 and below, 706 to 755, 756 to 779, 780 to 795, and 796 to 850. Likewise, the (self-reported) monthly income of the borrowers in our dataset is somewhat higher than the overall population, and we set five groups that partition the accounts about evenly: below \$3500, \$3500 to \$5000, \$5000 to \$7000, \$7000 to \$10,000, and over \$10,000 in monthly income.

The empirical distribution of credit utilization (as a percentage of total available credit) is highly concentrated at low values: over half of borrowers were using less than 20% of their available credit at the start of the deferred interest promotion. Dividing credit utilization percentages into equally sized groups would result in little meaningful variation among the lower four groups, so we instead set the boundaries to differentiate between substantially different levels. When conditioning on utilization, the five groups are: 0 to 10%, 10% to 20%, 20% to 40%, 40% to 70%, and 70% to 100%. Despite each successive band being wider, each group is successively smaller than the last.

A key indicator for whether a borrower is on track to successfully pay off the deferred interest debt during the promotional period is whether they pay *at least* one-twelfth of the original debt each period. An individual who makes such a payment in all twelve months will, by definition, succeed at the promotion; one who consistently pays below this threshold is likely to arrive in the final promotional month with substantial remaining debt. To home in on this distinction, we partition all non-terminal, non-missed payments[22] into five mutually exclusive and exhaustive groups. The terms of the promotion require borrowers to pay at least \$20 or 1% of the original debt each month, whichever is greater; we label all payments under \$30 as "trivial" or "minimal" payments.[23] Payments that exceed this minimal level but still fall short of the one-twelfth threshold (about 8.3% of original debt) are labeled as "below one-twelfth". The data shows a significant clustering of payments at *exactly* one-twelfth of the deferred interest debt (and just above), so we categorize payments from 8.3% to 9% of original debt as "near one-twelfth".[24] We label payments between 9% and 11% of the original debt as "above one-twelfth"; repeated payments in this range would pay off the full debt in ten or eleven months. Finally, payments in excess of 11% of the original debt are categorized as "high payments"– those that would complete the promotion in nine months or less, considerably less than the allowed time.

The partition of payments in the prior paragraph depended on single monthly payments; we also specify a related categorization of accounts depending on *cumulative* payments. In any month *after* the beginning of the promotion, a borrower is labeled as a "fast-payer" if the cumulative fraction of the debt they have already paid exceeds 9% times the number of prior months of the promotion (9% at the start of month 2, 18% at the start of month 3, etc). Such accounts are cumulatively on track to complete the promotion in *less than* twelve months. If instead the cumulative fraction of debt already paid is less than 8% times the number of prior months (8% at the start of month 2, 16% at the start of month 3, etc), the borrower is a "slow-payer" and is not on track to finish the promotion in twelve months. The middle cumulative payment group ("mid-payers") are on schedule to complete the promotion in exactly twelve months, if they persist in their average payment speed.

5.4 Moment Selection and Identification

A model that seeks to reproduce complex behavioral patterns can potentially be identified on a wide variety of data features– the number of potential data moments vastly

[22]Missed payments and payments that terminate the promotion by paying off the remaining debt are handled by other moments.

[23]The strong majority of such payments are exactly \$20.

[24]Consistent payments of 9.1% or more would pay off the original debt in eleven months or less.

exceeds the number of parameters to be identified. We have chosen a subset of these data features that reflect the qualitatively most important repayment behaviors of deferred interest borrowers, described in the paragraphs below. Each category of moments is designed to focus on identifying variation in the data: behavior that reveals the distribution of consumer preferences and beliefs as they vary with observable differences in borrowers. [25]

5.4.1 Heterogeneity in Missed Payment Probability

The α parameters govern simulated borrowers' probability of missing a monthly payment, ς_i, depending on observed characteristics and unobserved heterogeneity. These parameters can be identified fairly directly based on the observed rates of missed payments in the data, conditional on observables. For each observable group with respect to age, credit score, monthly income, and credit utilization, we calculate the fraction of months in which the borrower makes no payment at all (but still has deferred interest debt remaining). The overall average missed payment rate identifies α_0, while the coefficients α_j, α_s, α_y, and α_c are respectively identified by the slope of the missed payment rate across observable categories.

In the absence of unobserved heterogeneity in the probability of missing a monthly payment, the four observed characteristics would strictly determine each borrower's miss probability. The model would thus make strong predictions about the distribution of the number of missed payments across accounts, conditional on (say) credit score. More realistically, these four simple characteristics are not fully determinative of the missed payment rate, and the model's conditional distribution of missed payments would not match the data. To identify σ_α, the extent of unobserved heterogeneity in the missed payment rate, we include the distribution of the number of missed payments across borrowers conditional on credit score group. The *more similar* these distributions are across credit score groups, the *more unobserved heterogeneity* there must be.

For each month of the promotion, we compute the fraction of active borrowers in each of the five payment categories conditional on each observable group. That is, one moment represents the fraction of active borrowers aged 43 to 52 who pay between 9%

[25] All of the moments are calculated on the set of borrowers who remain in the promotion after the first month; we exclude those who exit immediately. There are very few such accounts in the empirical data, but many in the simulated data. If the model were specified slightly differently, these are the agents who would *never have entered* the promotion– it looked unappealing to them from the start.

and 11% of their original debt in month 7. With twelve months, nineteen observable groups, and five payment categories, there are 1140 moments of this variety.[26]

5.4.2 Obliviously Overconfident Borrowers

The distribution of payments in the early months of the promotion identifies the distribution of borrowers' *beliefs* about their probability of missing a monthly payment and whether they will be able to notice the final month of the promotion (and thus pay any remaining debt). To illustrate this idea, we will present a series of hypothetical borrowers with different beliefs about themselves, and discuss that borrower's optimal repayment heuristic (temporarily setting aside idiosyncratic preferences). Differences in the empirical distribution of (categorical) payments across observable groups thus reveals how borrower beliefs vary with those characteristics.

First, consider a borrower who believes they have an arbitrarily low probability of missing a monthly payment (say, one in ten thousand) *and* believes that they will notice the end of the promotion when it arrives. This individual believes that they can take full advantage of the deferred interest promotion by making minimal payments for the first eleven months, then paying all remaining debt in the final month– putting off repayment as long as possible. Across observable groups, the relative proportion of borrowers who make minimal payments thus identifies the distribution of *very confident* borrowers.

More generally, *any* borrower who believes that they would not notice the looming end of the promotion would want to choose a repayment plan that finishes within twelve months, without relying on a larger final payment. An individual who believes they are substantially likely to miss one or more monthly payments would optimally choose to pay *more* than one-twelfth each month, to provide a buffer against potential errors, but the one-twelfth threshold is a firm lower bound on their optimal payment heuristic. The combined proportion of accounts that pay near one-twelfth, above one-twelfth, or make a high payment (conditional on observable group) identifies the distribution of borrowers who believe they would not notice the end of the promotion– the ν parameters.

5.4.3 Self-Aware Overconfident Borrowers

Next, consider a borrower who expects that they will never miss any monthly payments, but believes that they would *not* notice when it's the final month of the promotion.

[26] By construction, the month 12 moments only include accounts who *fail* the promotion, as they made a non-terminal payment in the final month.

Such a borrower is confident that they will make any planned payment, but has to set a plan that pays off "naturally" within the promotional period, without relying on making a different payment in the final month. Under the modeling assumption of repeated constant payments, this individual's optimal plan would be to pay the deferred interest debt in twelve equal installments, finishing right on time. The deferred interest data exhibits significant clustering of payments very close to one-twelfth of the original debt, with significant variation in the proportion of such payments by credit score and utilization. The relative likelihood of paying about one-twelfth of the debt identifies the distribution of borrowers who are confident about making payments, but *not* changing payments in the final month.

Among borrowers who pay at least one-twelfth of the original debt each month, there is considerable variation in payment size. In the first two months of the promotion, *about half* of accounts that pay at least one-twelfth of the debt make "high payments" of over 11%; this fraction naturally falls over time as the fastest payers complete their repayment and exit. The proportion of borrowers who hew very close to paying one-twelfth are more likely to believe they won't miss any payments, while those who pay more than one-twelfth are more likely to believe they might miss a monthly payment. The relative frequency of "near one-twelfth" versus "above one-twelfth" payments across observable groups thus identifies variation in the distribution of perceived miss probability with those characteristics– the β parameters.

5.4.4 Sub-Extreme Payments

Now consider a borrower who is fully confident that they will notice the end of the promotion and pay any remaining debt in month 12, but believes that they might forget to make one or more monthly payments– including potentially the critical final month. For such a borrower, a missed monthly payment will cause them to incur a fee, but *does not* meaningfully increase their risk of facing a large deferred interest charge, as long as they make their month 12 payment. As long as the believed miss probability is not too high, it would still be optimal for this borrower to make minimal payments for the first eleven months. That is, making payments above the minimum level would reduce expected deferred interest paid (because less deferred interest is accrued when more of the debt is repaid), but *only* by the (relatively small) probability of a missed payment; this reduced cost is unlikely to overcome the smaller benefits of delaying repayment.

In our dataset, only about one-third of payments that are under the one-twelfth threshold are categorized as minimal payments ($30 or less). A significant fraction of borrowers make payments that do not put them on track to succeed at the promotion, but don't take full advantage of the promotion; how does the model account for such behavior? The model has (at least) two channels to account for "below one-twelfth" payments. First, the structural model includes the ω parameter as a scaling factor for a quadratic penalty on larger-than-usual final payments. The magnitude of this penalty is largely identified by borrowers' unwillingness to pay only the minimum required. Second, our model includes preference shocks for each repayment heuristic, idiosyncratic to each borrower. That is, some borrowers choose to pay \$50 or \$100 each month (rather than \$20) merely because they want to, for reasons *outside the model*– a pure structural error. Such borrowers are "leaving money on the table" so to speak, so their propensity to do so (relative to the making the minimum payment) helps identify σ_η, the scale of idiosyncratic preference shocks across plans.

Likewise, about one third of deferred interest borrowers begin the promotion by making payments in excess of 11% of their original debt. In the absence of idiosyncratic preference shocks, such repayment behavior would only be optimal if the borrower believed they were likely to miss *multiple* monthly payments. The data shows a non-trivial fraction of borrowers who appear to plan to fully repay the promotional debt in three to eight months. Rather than try to explain such behavior with simulated model agents that believe there is (say) a 25% to 75% chance they will miss each monthly payment, we instead simply include idiosyncratic preference shocks to reproduce these choices, providing further identification of σ_η.

5.4.5 Changes in Repayment Heuristics

Most payments in the deferred interest dataset are repetitious: the borrower makes the same payment each month. However, borrowers do occasionally seem to change repayment heuristics, switching to a faster or slower strategy– or suddenly exiting the promotion altogether by repaying all remaining debt. In the model, borrowers choose a repayment plan and believe they are making a once-and-for-all decision; in fact, they are randomly granted the chance to change their heuristic with probability π. Suppose π were large, so that borrowers *could* change their heuristic frequently. The model would predict that borrowers would *increase* the size of their payments over the course of the promotion. That is, as months pass, the time horizon over which to gain the benefits of delayed repayment *decreases*, while the potential deferred interest penalty *increases* (due to DI already accrued). A borrower given the chance to choose a new constant repayment plan would prefer to repay *faster* than they originally chose (on average). In contrast, the data show essentially no increase in the size of repayments over the course of the promotion. The absence of a secular upward trend in repayment size identifies the reoptimization probability π, which will be estimated to be relatively small– infrequent opportunities to change heuristics.

To the extent that borrowers change repayment behavior during the promotion, and that these trends are (slightly) different across observable groups, these data patterns identify the "drift" parameters δ. While the ν parameters govern whether model agents *begin* the promotion believing that they will (not) notice the final month of the promotion, these beliefs are not fixed: each month, active borrowers' latent belief level is shifted by δx_i. Depending on the sign of this drift rate and the initial level of latent beliefs at $t = 0$, the agent's perception of whether they will notice the arrival of the final month can flip if the latent belief changes sign. Such borrowers will not (necessarily) get a chance to change their repayment behavior immediately, but are likely to want to make a substantial change in their heuristic when given the opportunity (as motivated above). Differences across observable groups in *changes over time* in the payment category distribution thus identifies the belief drift parameters.

5.4.6 Rounding Preferences

More directly, the model includes the rounding preference shifters κ_{50} and κ_{100} to reflect the notion that making a simple round payment is less mentally taxing than trying to divide the debt among several periods. The moments of the objective function include the share of payments that are rounded to \$100 in each month of the promotion, as well as the share that are rounded to \$50 (but not \$100). Fixing the level of other the other parameters, these moments straightforwardly identify the rounding shifters.

5.4.7 Correlation in Unobserved Heterogeneity

The most important outcome for borrowers is whether they successfully pay off the original debt within the promotional period. We include as moments the fraction of accounts that succeed at the promotion conditional on observable group and the number of missed monthly payments (zero, one, or multiple). The level of the success rate, and its slope with respect to borrower characteristics, are "top line" measures of whether the model is fitting the key outcome. Differences in the success rate by number of missed payments *within* each observable group identifies the extent to which the unobserved factors that contribute to an individual's miss probability are correlated with the unobserved factors that determine whether they will "critically fail" by not noticing the final month of the promotion– the correlation coefficient ρ_A.

To see this, suppose these unobservable components were fully independent, and consider borrowers in one observable group– say, those with credit scores below 706. The success rate for accounts who miss one payment would be lower than for those who miss none (and multiple missed payments would be even lower), due to correlation with *other* observables, but these differences would be slight or moderate. That is, a low credit score person who missed one payment is more likely to have *other* characteristics associated with missing payments (like being younger), and those characteristics are generally also predictive of failing the promotion. The data, however, shows that the number of missed payments is a *very strong* predictor of successfully paying off the debt in time– much stronger than would be the case if unobserved heterogeneity in miss probability and critical failure probability were uncorrelated. This feature of the data thus identifies the extent of correlation between the two unobserved factors.

5.4.8 Early Exit from the Promotion

Beyond the overall success rate, we also care about *when* borrowers complete their repayment, as a significant share of accounts pay off their debt well before the deadline. We define the monthly exit rate as the proportion of active accounts that pay off all remaining deferred interest debt in that month. In the data, the exit rate is low but slowly increasing in the second through eighth months of the promotion, accelerates slightly in the ninth and tenth months, then jumps to a 40-50% exit rate in the final two months. In the estimation, we jointly condition the exit rate by observable group, time in the promotion (months 3-9, month 10, and months 11-12), and the cumulative fraction of the debt paid (slow-payers, mid-payers, and fast-payers, described above).

Among the borrowers who exit the promotion well before it ends– say, before month 10– some appear to have originally planned to pay off quickly (e.g. paying \$400 per month on a \$2500 debt). However, a significant number exit *suddenly*, deviating from a steady stream of ordinary payments by making a large terminal payment. In the structural model, we interpret these borrowers as those who were given an unexpected opportunity to choose a new payment heuristic, and who realized that their continued participation is *no longer worth it*. Formulating a payment plan and making monthly payments is more mentally taxing than paying off the deferred interest debt; the model captures this with the τ_1 preference shifter: a bonus to utility from paying all remaining debt. The rate at which mid-payers and slow-payers exit the promotion in months 3-9 thus identifies this cost of participation parameter: the promotion was attractive in the first month, but after much of the “delay value” of putting off repayment has been captured in the early months, the ongoing effort of making payments isn’t worth it anymore.

5.4.9 Averting Critical Failure

The deferred interest data shows that while cumulative prior payments are predictive of exiting the promotion (i.e. a borrower who has already paid more of the original debt is more likely to finish paying their debt in a given month), these exit rates do not vary much across observable groups in the first ten months of the promotion. In the final two months of the promotion, both mid-payers and fast-payers are on track to pay off the debt by month 12 simply by making their “typical” payment; the exit rate in these months likewise does not vary with borrower characteristics. However, in order for slow-payers to succeed late in the promotion, they *must* actually notice that month 12 has arrived, and the data reveal that borrower characteristics *are* predictive of the exit rate. The μ parameters that govern whether borrowers *actually* notice it’s month 12 (versus critically failing by treating it like an ordinary month) are thus identified by the exit rate in months 11 and 12 for slow-payers– those for whom failing to notice it’s month 12 *actually matters.*

6 Results

This section begins by discussing the results of the structural estimation, including the model’s fit to the data and interpretation of the estimated values. It then proceeds to an analysis of the consumer welfare effects of deferred interest promotions by decomposing the impact of various model features on financial outcomes.

6.1 Estimated Model

The estimated structural parameters and their standard errors are presented in Table X. The estimation includes 1431 moments to be matched while estimating 35 parameters– a significantly overidentified system. As a basic quantitative measure of model fit, the objective function reaches a minimum value around 62,646.5; as is typical with complex structural models, this does not come close to satisfying the overidentification J-test. [27]

[27] If the model were true, the minimum of the objective function should be distributed χ^2 with $M - K = 1396$ degrees of freedom, with respect to the empirical sample.

Table 2 Parameters Estimated by the Simulated Method of Moments

Param	Description	Value	Std err
α_0	Actual missed payment probability: constant term	-0.597	(0.016)
α_j	Actual missed payment probability: coefficient on age	1.63e-4	(4.40e-4)
α_s	Actual missed payment probability: coefficient on FICO score	3.42e-3	(0.07e-3)
α_y	Actual missed payment probability: coefficient on log income	7.89e-2	(0.61e-2)
α_z	Actual missed payment probability: coefficient on credit utilization	-4.31e-3	(0.17e-3)
σ_α	Actual missed payment probability: stdev of error term	0.930	(0.007)
μ_0	Actual critical failure probit: constant term	-1.844	(0.052)
μ_j	Actual critical failure probit: coefficient on age	-3.50e-3	(1.10e-3)
μ_s	Actual critical failure probit: coefficient on FICO score	1.98e-4	(1.85e-4)
μ_y	Actual critical failure probit: coefficient on log income	0.329	(0.018)
μ_z	Actual critical failure probit: coefficient on credit utilization	-2.22e-2	(0.05e-2)
β_0	Believed missed payment probability: constant term	7.485	(0.037)
β_j	Believed missed payment probability: coefficient on age	-7.12e-2	(0.18e-2)
β_s	Believed missed payment probability: coefficient on FICO score	-2.32e-2	(0.02e-2)
β_y	Believed missed payment probability: coefficient on log income	2.044	(0.020)
β_z	Believed missed payment probability: coefficient on credit utilization	-2.04e-2	(0.05e-2)
σ_β	Believed missed payment probability: stdev of error term	10.63	(0.04)
ν_0	Believed critical failure probit: constant term	4.953	(0.014)
ν_j	Believed critical failure probit: coefficient on age	-2.26e-2	(0.04e-2)
ν_s	Believed critical failure probit: coefficient on FICO score	-6.44e-3	(0.07e-3)
ν_y	Believed critical failure probit: coefficient on log income	4.23e-2	(0.59e-2)
ν_z	Believed critical failure probit: coefficient on credit utilization	1.50e-2	(0.02e-2)
δ_0	Drift in belief in critical failure probit: constant term	-3.50e-2	(0.09e-2)
δ_j	Drift in belief in critical failure probit: coefficient on age	2.87e-3	(0.29e-3)
δ_s	Drift in belief in critical failure probit: coefficient on FICO score	6.59e-4	(0.47e-4)
δ_y	Drift in belief in critical failure probit: coefficient on log income	-8.21e-2	(0.44e-2)
δ_z	Drift in belief in critical failure probit: coefficient on credit utilization	5.73e-4	(1.01e-4)
ρ_A	Correlation b/w error terms for actual mistakes	0.602	(0.009)
ρ_B	Correlation b/w error terms for perceived mistakes	0.180	(0.008)
π	Probability of getting to choose new plan	9.30e-2	(0.04e-2)
σ_η	Scale of preference shocks over payment plans ($)	25.12	(0.05)
σ_η	Scale of preference shocks over payment plans ($)	25.12	(0.05)
ω	Magnitude of penalty for large final payment	8.31e-6	(0.04e-6)
κ_{50}	Preference bonus: payment rounded to $50	-11.13	(0.06)
κ_{100}	Preference bonus: payment rounded to $100	16.01	(0.06)
τ_1	Preference bonus: paying all remaining debt	33.34	(0.17)

Qualitatively, however, fit of the model is excellent, with essentially all targeted features being matched in level, slope, and distribution (as appropriate). Figures showing the graphical fit of all of the estimation moments are presented in the appendix.

The estimated values of the α parameters governing the actual probability of missed payments indicate that all of the observed characteristics (except age) are significantly predictive of missed payments. Unsurprisingly, borrowers with higher FICO scores and/or lower credit utilization are less likely to miss a monthly payment, as are higher income borrowers (to a lesser extent). The estimated value of σ_α generates a correlation coefficient of $r \approx 0.23$ between the distribution of $\alpha \cdot x_i + \epsilon_i^s$ and the base $\alpha \cdot x_i$. [28] That is, the four observed characteristics are predictive of miss probability, but there is significant unobserved heterogeneity.

In contrast, the β parameters for borrower *beliefs* about their probability of missing a payment reveal that estimated beliefs about miss probabilities are much different. The estimates of β_0 and σ_β indicate that model borrowers are, on average, much more confident about not missing payments, but that beliefs are also much more dispersed than actual probabilities. With such a high mean and wide spread of $\beta \cdot x_i + \xi_i^s$, the model generates large masses of simulated borrowers who believe they will *always* or *never* miss monthly payments– very low and very high values respectively. Model borrowers who believe they will almost always miss monthly payments will not participate in the promotion, paying all of the debt in the first period; such borrowers are not included in further moment calculations. Those who remain in the promotion are thus a censored slice of a normal distribution, indicating that the assumption of an underlying normal belief distribution might not be valid. Among participating borrowers, the majority believe they can (essentially) never miss payments– the kind of overconfidence that leads them to gambling on making a concluding payment in the final month of the promotion. Given the values of β_0 and σ_β, the coefficients on observable characteristics serve as shifters on the fraction of borrowers with plausibly realistic beliefs about missing a payment versus those who treat it as an impossibility.

Estimates of the μ parameters, concerning the realization of critical failures (not noticing the final month), yield no surprises. The data show significant variation in success rates by demographics for accounts that reach the final two months of the promotion as slow-payers (under 8% of original debt paid per month on average), and the model matches these rates by adjusting μ. Recalling that higher values of $\mu \cdot x_i$ are associated with lower critical failure propensity, the estimation finds that higher

[28]For miss probability ς_i itself, $r \approx 0.18$ between versions with and without the error term.

income borrowers and those with lower credit utilization are significantly more likely to successfully complete repayment in month 12. Younger borrowers and those with higher credit scores are slightly less likely to critically fail, but the relationship is not strongly significant– the observed relationship is accounted for by correlation between credit score and utilization. Conditional on an observable category, success in the promotion is strongly predicted by the number of missed payments (zero, one, or multiple), yielding a strongly positive correlation in the error terms on actual miss probability and actual critical failure: $\rho_A \approx 0.6$.

Model borrower beliefs about whether they will be able to notice the arrival of the final month (the ν parameters) indicate significant heterogeneity by observable characteristics, with largely an *inverse* relationship with actual probabilities. Particularly, while high credit utilization is positively associated with actual critical failure, such borrowers *act as if* they will surely notice the end of the promotion– μ_c is estimated to be negative while ν_c is negative. Likewise, borrowers with higher credit scores repay their debt with seemingly more caution about the prospect of critical failure, choosing heuristics that fully repay the debt "naturally", even though they are actually *less* likely to fail in this way. Higher income borrowers are less likely to critically fail, but income has an essentially flat relationship with beliefs.

The model allows borrowers to adjust their binary belief about critical failure over the course of the promotion, with the latent belief value drifting at rate $\delta \cdot x_i$ each month. Simulated borrowers whose latent belief changes sign during the promotion might want to pursue a drastically repayment plan, ensuring that they repay the debt by month 12 or relying on a large final payment. At the estimated parameters, all borrowers have a negative drift rate, making them *less confident* as the promotion proceeds and potentially coming around to the dire prospect of critical failure. The significantly negative value of δ_y means that higher income borrowers *learn faster* and are more likely to switch from being optimistic about critical failure to pessimistic. Conversely, older borrowers and those with higher credit scores or credit utilization have less negative drift rates.

However, the interpretation of these coefficients is somewhat ambiguous: those with high credit scores might adjust their beliefs more slowly because their initial beliefs were *more realistic*– they have less to learn about themselves. Note that *all* borrowers have a substantial likelihood of failing to notice the final month of the promotion– even the most capable groups make this critical mistake 20% of the time. A borrower with "correct" beliefs recognizes that failing to notice the final month is a realistic possibility, and hence be fully pessimistic under our specification of binary perceptions.

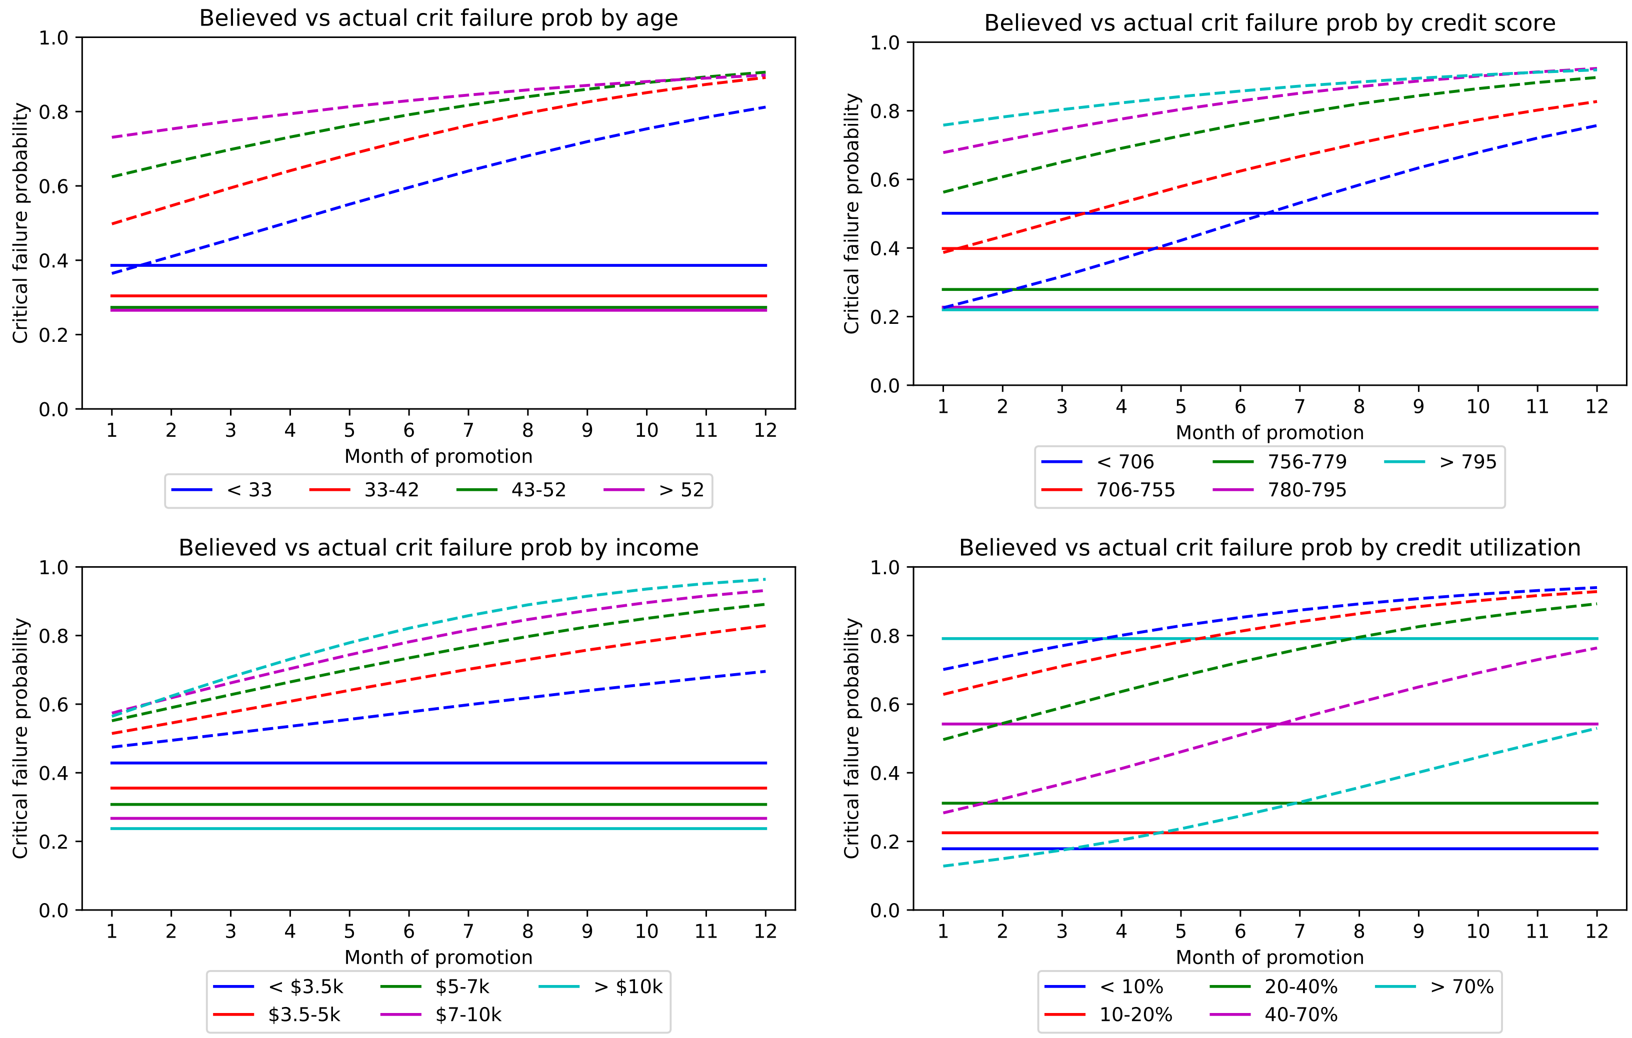

Figure 6 Actual (solid) and perceived (dashed) fraction of borrowers in each observable group who will fail to notice the final month of the promotion.

The estimated model finds that borrowers have about a 9% probability each month to be given the opportunity to choose a new repayment heuristic, contrary to their expectation of a once-and-for-all decision at the start of the promotion. On average, model borrowers will get one such chance during the yearlong promotion, consistent with the steady repayments seen in the data. When they select a plan, the model judges that borrowers exhibit significant tendencies to make payments rounded to \$100 (yielding effective utility worth about \$16 relative to a timing-based heuristic). The estimated model reveals that participation in the promotion is taxing, so borrowers will not participate unless their subjective expected return exceeds $\tau_1 = \$33.34$. This value is reasonable when interpreted as the anticipated effort cost of making payments each month: about \$3 per payment when spread over ten to twelve months in the promotion.

Moreover, model borrowers are somewhat reluctant to plan to make final payment significantly larger than their ordinary payment. To interpret the estimated magnitude of $\omega = 8.31 \times 10^{-6}$, we use the example of a borrower with initial debt of \$1500 who is supremely confident in their ability to neither miss payments nor fail to notice the end of the promotion. This borrower would likely want to delay repayment to take full advantage of the promotion, and thus primarily considers the slowest paying heuristics: \$100 per month, \$50 per month, and the minimum payment of \$20. Under these three plans, the borrower expects to pay respectively \$300, \$900, and \$1260; at the estimated ω, these would generate utility penalties of \$0.75, \$6.73, and \$13.19. Given these penalties and idiosyncratic preferences (discussed below), some borrowers would opt to repay slightly faster than the minimum rate, but still likely below the critical one-twelfth threshold given their level of confidence.

Finally, we find that there is a significant idiosyncratic component to borrower choice, unaccounted for by model features: the standard deviation of preference shocks over heuristics is $\sigma_\eta = \$25.12$. With an average initial debt of around \$1500 and a 5% preference rate for delaying repayment, taking full advantage of the promotion by making *no payments* until the final month (ignoring the minimum payment for simplicity) yields a "delay value" around \$75 (relative to paying the entire debt immediately). Differences in the subjective value of similar repayment plans are on the order of a few dollars, so unexplained idiosyncratic preferences drive the fine-tuned decisions of borrowers. That is, the model can only explain large qualitative differences in repayment behavior across the population, not specific borrower choices.

6.2 Counterfactual Welfare Analysis

We now turn attention to an analysis of how borrowers' several deviations from or ideal repayment behavior impacts them financially. Recall that a hypothetical "perfect borrower" who has no human frailties or biases, cannot make any kind of mistake, and is fully rational and cognizant of these facts would take full advantage of the deferred interest promotion by making the minimum required payment until the final month, when all remaining debt is repaid.

In contrast, model borrowers (whose behavior closely replicates that of actual borrowers) do not conduct themselves as would a perfect borrower for three primary reasons. First, borrowers can *make mistakes* during the promotion, both by missing monthly payments and by failing to notice the end of the promotion. Second, even though they recognize that such mistakes are possible, borrowers have *incorrect beliefs* about the probabilities of such events. Third, borrowers have *cognitive biases* and *idiosyncracies* that lead them to make round payments or choose a plan for non-financial reasons outside the model. In this section, we counterfactually simulate borrowers' behavior when each of these channels is turned off in sequence, transforming estimated model borrower behavior into perfect borrower behavior step-by-step.

Our primary outcome of interest for this exercise is the *financial benefit* of the deferred interest promotion to the borrower. This measure includes the borrower's value of delaying repayment of the debt (at the calibrated annual rate of 5%) as well as the costs of missed payment fees and deferred interest (for those who fail the promotion), all in expectation. It does *not* include non-financial components that affect the borrower's decision, including utility bonuses for round payments and the preference shock for the chosen repayment plan. For each counterfactual scenario, we re-solve the deferred interest model, then simulate repayment histories for one hundred copies of each borrower in our estimation data sample and compute average financial benefits conditional on each observable group.

The results of these experiments are presented below in Figure 7; in each figure, the blue curve represents borrower outcomes under the estimated model, labeled as the baseline. On average, borrowers derive rather modest financial benefits from participating in the deferred interest promotion, about $8, or one half of one percent of the average purchase size. There is some heterogeneity in realized benefits under the baseline, as the riskiest borrowers (those with low credit scores and/or high credit utilization) have net *negative* financial flows from the promotion, while more able borrowers experience

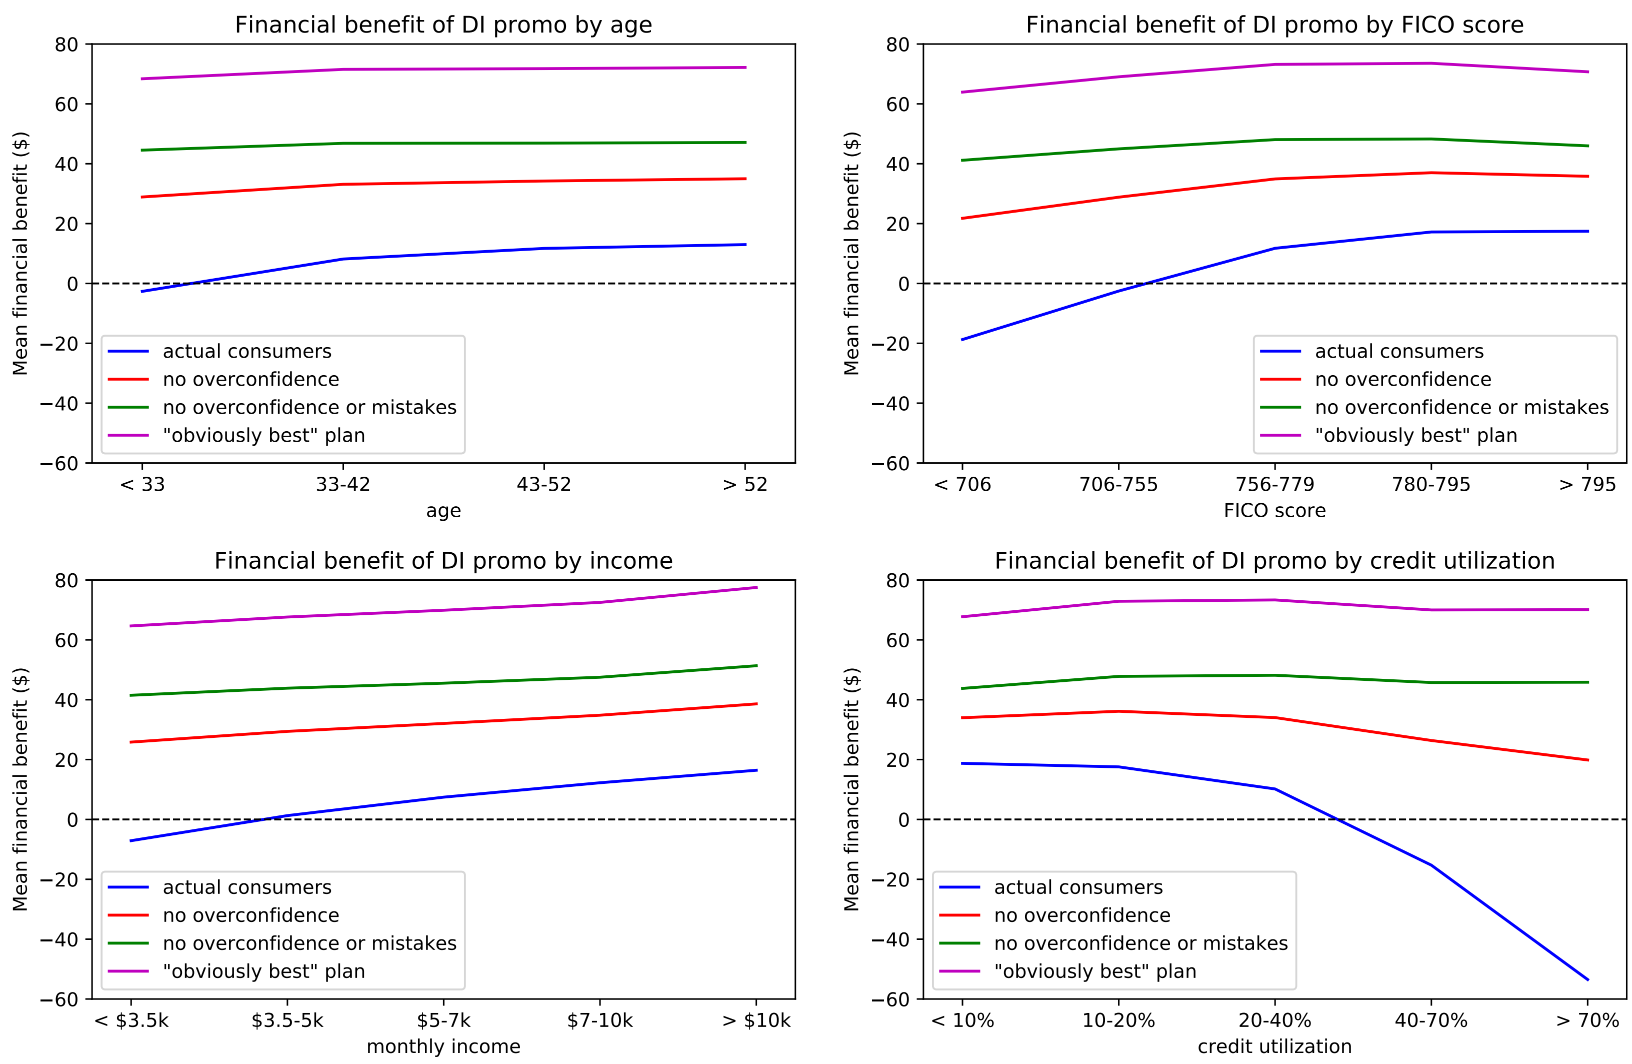

Figure 7 Financial benefits of participation in the deferred interest promotion by observable group under different counterfactual specifications.

consistently positive benefits due to both fewer late fees and paying less deferred interest. Note that if the preference rate for delaying repayment had been calibrated solely on potential returns to short run saving (say, 1 to 2%), then the calculated financial benefits in the baseline would likely be negative for most borrowers, as the potential gains from the promotion would be so much smaller, but the financial risks the same.

For our first counterfactual experiment, we consider what would happen if borrowers could never make mistakes while repaying their debt– neither missing payments nor failing to notice the end of the promotion. In the model, this is achieved by setting α_0 and μ_0 to arbitrarily high values, so that ς_i and φ_i are both zero for all borrowers, per (1). Such counterfactual borrowers still *believe* they might make mistakes and choose their payment heuristic accordingly; in fact, they will behave exactly as they would in the baseline. The financial benefits accruing to consumers in this situation are shown on the green trends in the figures. With any chance of failure removed, borrowers reap significantly greater benefits, which are much flatter across demographic groups (relative to the baseline), ranging from about \$33 to \$36. Indeed, the most capable borrowers (with high credit scores and low credit utilization) actually attain slightly *lower* financial benefits than riskier borrowers, precisely because they have more cautious *beliefs* about their propensity for errors, and hence leave more money on the table by employing more conservative strategies.

Consider instead an environment in which borrowers might miss payments or fail to notice the impending deadline for the promotion, but they correctly understand these probabilities rather than having errant beliefs. That is, for each model agent, we set $\widetilde{\varsigma}_i$ to the true value of ς_i and grant the agent foresight about whether they will critically fail: $\widetilde{\varphi}_i = \mathbf{1}(\mu \cdot x_i + \epsilon_i^f < 0)$. Because these beliefs are correct, we also assume that borrowers do not learn about critical failure during the promotion, setting all δ parameters to zero. With these changes, model borrowers are able to make more informed decisions about how to repay their deferred interest debt. Borrowers who are highly unlikely to miss payments nor critically fail at the end of the promotion will be much more likely to take full advantage of the terms of the deal, delaying repayment as long as possible. Conversely, borrowers who are riskier will repay the debt more cautiously, accounting for the possibility of missing payments and the fact that they will not be able to make a larger than usual final payment. Consequently, borrowers of all stripes will (on average) experience lower costs and greater benefits with this insight about themselves, as shown on the red trends of Figure 7. For most observable groups, the financial benefits from having correct beliefs are very close to those from never making mistakes. However, while the riskiest borrowers (those with low credit scores or high credit utilization) gain the most from their newfound knowledge, as they make the largest errors in the baseline, the financial benefits are not as large as if they could not make mistakes– they need to act conservatively when repaying the debt so as not to risk failing the promotion! Hence the red trends are somewhat increasing in borrower quality, whereas the green trends are very slightly decreasing in those dimensions.

Borrowers who can't actually make mistakes but act as if they can are effectively leaving money on the table: they could take complete advantage of the deferred interest promotion, but don't know it. Combining the first two counterfactual experiments, the magenta trends plot average financial benefits if borrowers can neither miss payments nor fail to notice the end of the promotion *and* they are aware of this fact. The level of these curves can be interpreted in two ways, relative to the green and red trends from the prior experiments. The gap between the magenta and green trends represents the value of borrowers *knowing* they can't make mistakes, so that they are much more willing to delay repayment. This benefit is slightly larger for more capable borrowers, who have more conservative beliefs in the baseline. Alternatively, the vertical gap between the magenta and red trends represents the benefits of *not making mistakes* even when borrowers have a perfect understanding of themselves. As expected, this benefit is largest for the riskiest borrowers, as they have the most margin for improvement with respect to repayment mistakes.

Recall that the average initial debt is slightly over \$1500, so that fully postponing repayment for one year would yield a delay value around \$75. However, counterfactual borrowers who know that they can't make mistakes only reap financial benefits around \$46 on average. The difference between the simple back-of-the envelope calculation and the height of the magenta trend is accounted for by the rounding biases and idiosyncratic preferences of model borrowers. That is, even with knowledge of their ability to execute any repayment plan perfectly, the prior experiment still allowed borrowers to prefer making round payments *and* to have preference shocks over the various repayment heuristics. Additionally turning off these features results in the hypothetical "perfect borrowers" described at the top of this section, who will always choose to make minimum payments for first eleven months and then repay the remainder at the end of the promotion. Consequently, the cyan "perfect borrowers" trend is nearly flat just below \$75, as predicted; any variation in this trend is driven by differences in average purchase size across observable groups (e.g. higher income borrowers make larger purchases). Relative to the magenta trend, perfect borrowers benefit almost uniformly across observables, as they are all correcting the same extent of error from their trembling hands. The degree of human idiosyncracy that we estimate has a significant financial impact on borrowers: between 20% (for the riskiest) and 45% (for the most capable) of the difference in financial benefits between perfect borrowers and the baseline model is accounted for by these features.

7 Conclusion

It is not controversial to suggest that people are aware that they have cognitive limitations, or that self-aware people (or organizations) construct "choice architectures" designed to guard against those limitations (cf Simon (1950); Sims (2003); Laibson (1997)). What is more difficult is to measure how effective such prudential (Kimball (1993)) efforts may be in guarding against the consequenses of those limitations. Even prudent and self-aware consumers have no access to systematic data on the prevalence of the kinds of errors they seem to guard against, so consumers' perceptions of their own limitations might be systematically biased.

This paper provides some evidence from a rare case in which people's measurable behavior reveals the degree to which consumers are self-aware. We specify a structural model of borrowers that are qualitatively slef-aware but may be quantitatively overconfident of their abilities. Our model borrowers choose persistent repayment heuristics to manage their chances of failing the promotion, but some do so under the critically erroneous belief that they will surely pay any remaining debt at the end of the promotion. In fact, all kinds of borrowers in our data are substantially likely to blunder through the promotion's deadline, making their typical monthly payment rather than repaying all their debt. Our estimation reveals that the observable characteristics most predictive of repayment mistakes are also associated with overconfidence, i.e. the least capable borrowers are also the least sophisticated in their self-knowledge. Counterfactual analysis of borrower behavior reveals that most borrowers would be about equally improved by either gaining insight into their own probability of making mistakes or being unable to make mistakes at all. However, while the least savvy borrowers would gain the most from being fully self-aware, they would benefit even more from not making mistakes at all.

We leave it to future research to consider the implications for regulation, optimal market design and other questions.

References

Consumer Financial Protection Bureau (2013): "CARD Act Report," .

——— (2015): "CARD Act Report," .

Kimball, Miles S. (1993): "Standard Risk Aversion," *Econometrica*, 61(3), 589–611.

Laibson, David (1997): "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, CXII(2), 443–477.

Nelson, Scott Thomas (2017): "Private information and price regulation In the US credit card market," Ph.D. thesis, Massachusetts Institute of Technology.

Simon, Herbert Alexander (1950): *Administrative behavior.* Macmillan New York.

Sims, Christopher (2003): "Implications of Rational Inattention," *Journal of Monetary Economics*, 50(3), 665–690, available at http://ideas.repec.org/a/eee/moneco/v50y2003i3p665-690.html.