Explaining the Credit Card Debt Puzzle with Heterogeneous Liquidity Fungibility

Nick Pretnar *[1]

[1]*Carnegie Mellon University, Tepper School of Business*

March 17, 2019

1 Introduction

Why do some individual consumers simultaneously carry interest-bearing consumer debt, from say credit cards and student loans, while simultaneously saving? How do consumers who are more likely to engage in debt expenditure respond to exogenous income, wealth, and credit shocks? What are the distributional implications of this phenomenon? Depending on their preferences for engaging in debt expenditure, how does the relative position of consumers within the wealth and consumption distributions change in response to both idiosyncratic and aggregate shocks? How do aggregate shocks affect the summary statistics of the consumption and wealth distributions when we allow for cross-sectional consumer heterogeneity for debt utilization? Finally, how are neighborhood economic and demographic characteristics correlated with consumers' credit utilization preferences and susceptibility to adverse outcomes results from exogenous shocks? This working paper presents a model with mental accounting features designed to answer all of these questions. Future work will take the model to a unique dataset featuring debit and credit card expenditures at the transaction-level linked consumer by consumer.

The main theoretical contribution of this working paper is to answer the so-called "credit card debt puzzle" (Telyukova and Wright 2008). Telyukova and Wright (2008) call the credit card debt puzzle just "another manifestation of the venerable *rate of return dominance puzzle* from monetary economics." The authors argue that households need liquidity for purchases when using a credit card is difficult, similar to the argument made in Telyukova (2013). Writing in 2008, they note that riding in a taxi or buying a pack of cigarettes at a mom-and-pop bodega are generally cash-only transactions. Also, things like rent and mortgage payments cannot typically be made with a credit card. In today's world, now some 11 years later, the number of consumption items for which this

*Carnegie Mellon University – Tepper School of Business, 5000 Forbes Ave., Pittsburgh, PA 15218; npretnar@cmu.edu . Nick acknowledges support from the National Science Foundation Graduate Research Fellowship under Grant No. DGE1252522.

conjecture applies has greatly diminished. Payments for Ubers and Lyfts are automatically posted on credit card accounts. Every corner convenience store and gas station selling cigarettes takes a credit card.[1]

Gross and Souleles (2002) show that credit limit increases create a rise in debt, inconsistent with the PIH. The marginal propensity to consumer (MPC) out of liquidity (non credit) is highest for individuals near their credit limit. The authors note that previous studies which have dealt with how liquidity constraints distort consumption must often arbitrarily define how households are constrained, bases on some rule of thumb regarding net worth.[2] The set up presented here avoids such arbitrary modeling assumptions. Instead, the model allows for preference heterogeneity across consumers, directly capturing variation in credit utilization and consumption expenditure.

2 Preliminary Data Analysis

Given a dataset with linked debit and credit card accounts describing the transaction history for individual consumers, a preliminary analysis to either support or refute the work in Gross and Souleles (2002) will examine how individual MPC's vary between different liquidity sources, controlling for anticipated and unanticipated income shocks, unanticipated price shocks, unanticipated emergency expenditures, and recurring expenditures. Here, I describe regressions which can shine light on how MPC's fluctuate between consumers and across different liquidity sources. The proposed regressions are just extensions of those in Montgomery, Olivola, and Pretnar (2018) with more robust controls.

Let Δ be the backwards first-difference operator, $ch_{i,t}$ be available real checking balances, $m_{i,t}$ the credit limit, $b_{i,t}$ credit card debt, and $c_{i,t}(x_{i,t})$ be real consumption which depends on cash $ch_{i,t}$ or available credit $cr_{i,t} = (m_{i,t} - b_{i,t})$.[3] Let $MPC(x_{i,t})$ denote the marginal propensity to consume from the given liquidity source x, $l_{i,t}$ period t income, $\bar{l}_i$ mean income over the sample for agent i, $p_{i,t}$ an aggregate price level, $\overline{p}_{i,t}$ a period t moving average price level, $int_{i,t}$ the effective monthly interest rate, $emerg(x_{i,t})_{i,t}$ emergency expenditures, and $recur(x_{i,t})_{i,t}$ total recurring expenditures.[4] $u_{i,t}$ and $v_{i,t}$ are predetermined error terms. Seeking a preliminary assessment of this theory, I propose to use two-stage least squares (2SLS) regression to estimate the following simultaneous equations model using the proprietary bank data for which income can be imputed and emergency expenditure shocks identified under various assumptions:

$$\begin{aligned} MPC(ch_{i,t}) = \frac{\Delta c(ch_{i,t})_{i,t}}{\Delta ch_{i,t}} &= \beta_0 + \beta_1 \frac{l_{i,t} - \bar{l}_i}{\sigma_{l_{i,t}}} + \beta_2 l_{it} + \beta_3 \frac{p_{i,t} - \overline{p}_{i,t}}{\sigma_{p_{i,t}}} + \beta_4 \mathbf{1}[emerg_{i,t} > 0] + \eta_i^1 \\ &+ \gamma_1 \ln[emerg_{i,t}(ch_{i,t})] + \gamma_2 \ln[recur_{i,t}(ch_{i,t})] + \gamma_3 MPC(cr_{i,t}) + u_{i,t} \end{aligned} \tag{1}$$

[1]Exceptions to this trend are often met with suspicion, as "cash only" can be seen as a guise for something more nefarious given the ease with which the shop books can be cooked.

[2]Indeed, Browning and Lusardi (1996) and Gross and Souleles (2002) point out that the lack of consensus regarding the role of precautionary motives in forming liquidity constraints.

[3]Assume t subscripts are indexed by i, thus reading t_i.

[4]Note that *emerg* and *recur* depend on the liquidity source.

$$MPC(cr_{i,t}) = \frac{\Delta c(cr_{i,t})_{i,t}}{\Delta cr_{i,t}} = \theta_0 + \theta_1 \frac{l_{i,t} - \bar{l}_i}{\sigma_{l_{i,t}}} + \theta_2 l_{it} + \theta_3 \frac{p_{i,t} - \overline{p}_{i,t}}{\sigma_{p_{i,t}}} + \theta_4 \mathbf{1}[emerg_{i,t} > 0] + \theta_5\, int_{i,t} + \eta_i^2$$
$$+ \nu_1 \ \ln[emerg_{i,t}(cr_{i,t})] + \nu_2 \ \ln[recur_{i,t}(cr_{i,t})] + \nu_3\, MPC(ch_{i,t}) + v_{i,t} \tag{2}$$

In the above, the terms $\frac{l_{i,t}-\bar{l}_i}{\sigma_{l_{i,t}}}$ and $\frac{p_{i,t}-\overline{p}_{i,t}}{\sigma_{p_{i,t}}}$ are standardized deviations from mean income and moving average price expectations respectively. The modeling assumption is that marginal adjustments to consumption from credit and checking balances are endogenously determined, along with the right hand side expenditure variables $emerg_{i,t}(x_{i,t})$ and $recur_{i,t}(x_{i,t})$. After all, agents could choose not to engage in emergency expenditure and to pause or cancel recurring payments, thus affecting MPC's. We assume that the presence of an emergency expenditure is indeed exogenous, so that consumers must choose to positively spend in an emergency, but can decide the level of expenditure engagement and which method of payment to use, thus the term $\mathbf{1}[emerg_{i,t} > 0]$ does not depend on liquidity source. Heterogeneity is captured by η_i^1 and η_i^2, which are agent fixed effect terms.

The design of the regression can be used to understand the degree to which $MPC(x_{i,t})$ is sensitive to income and price shocks. One can test whether consumers are more likely to use credit cards for emergencies, i.e. $\theta_4 > \beta_4$. Violations of PIH[5] consistent with mental accounting theory would occur if $\theta_2 > \beta_2$, i.e. consumers increase consumption out of credit if income increases.

3 A Model with Mental Accounting

Suppose consumers behave according to the theory of Thaler (1985) whereby ex-ante they set expenditure plans for themselves according to some optimization rule. Such behavior would be consistent with what is known as the "planner's" side of a mental accounting problem. We assume consumers' plans are rational under mental accounting constraints. After planning, consumers go out and engage in expenditure — the "doer's" side of the problem. Under the planner's side of the problem, consumers choose budgets for how much they will spend using each separate method of payment — credit cards, cash, savings account liquidation, etc. Under the doer's side of the problem, consumers engage in expenditure, with deviations to ex-ante budgets following an exogenous, stochastic process. It should be noted that we use the term "planner" in a different context than is used in many economic settings. Here, the "planner" can be thought of as one side of the consumer's personality: he is the guy who sits down at his computer once a week, looking over bank statements trying to set a spending budget.[6] The "doer" on the other hand does not think about the future like the "planner" does. Rather, he simply engages in consumption expenditure in between the times the "planner" plans budgets. The "doer" possesses some weakly

[5]Permanent income hypothesis.

[6]This is not a "social planner" who knows the utility functions of all agents and chooses allocations so as to maximize some social welfare function. Still, the "planner" does indeed know his own utility function, and he does make planning decisions according to rational expectations.

latent recognition of the underlying budget set by the "planner," but his decisions are idiosyncratic and error prone. The underlying expenditure process thus contains a stochastic element that directly describes the degree to which the consumer misses his budget target.

An important feature of mental accounting will be exploited here to help explain the "credit card debt puzzle" described in Telyukova and Wright (2008) and Telyukova (2013). Namely, rather than forming a budget strictly out of total available liquidity which includes both unsecured credit plus available assets, individuals decide how much to spend with each separate source of liquidity. The formulation presented here thus admits a slightly modified definition of traditional Thalerian "fungibility." Whereas under "perfect fungibility" in the sense of Thaler (1985), consumers do not care where liquidity comes from, the formulation outlined below says that a consumer engaging in what I am calling "rational fungibility" will seek to equate the marginal utility of having more real assets next period with the negative marginal disutility of having debt. Thus, consumers face a tradeoff between debt and saving which neither results from precautionary liquidity holding motives as in Telyukova and Wright (2008) and Telyukova (2013) nor complicated hyperbolic time preferences as in Laibson, Repetto, and Tobacman (2000). Rather, the tradeoff is a specific by-product of the mental accounting structure: consumers decide how much money to spend on their credit cards, how much to spend in cash, and how much to save, separately. In the model, this amounts to splitting the budget constraint, thus allowing the stock of wealth and the level of debt to evolve according to separate laws of motion.

3.1 Model Outline

Time is discrete and indexed by t. A consumer has preferences $u(c_t)$ over consumption c_t. Consumers have three types of liquidity out of which they can consume — invested assets a_t, available cash z_t, and unsecured credit $m_t - b_t$, where m_t describes the maximum credit line and b_t is the total amount borrowed.[7] Note that consumers may be able to invest in multiple assets each with different returns and have multiple lines of unsecured credit each charging different interest rates. To simplify the illustration, suppose consumers can invest in only a single asset and use only a single line of unsecured credit. Each period the "planner" knows balances of cash z_t and available credit $m_t - b_t$. Consumers make payments on accrued debt according to e_t which is > 0 when paying down debt. Debt accrues gross interest Q_t at the beginning of each period. Certain kinds of debt instruments provide both residual benefits for usage — think travel points or cash back from a credit card — while also penalizing consumers for late payments or payments below some minimum threshold. These residual components are captured by the function $\zeta_t(e_t, \cdot)$, which is > 0 if a penalty is assessed.[8] Invested assets earn gross return $R_t > 0$, but which is < 1 if the investment depreciated and > 1 if it appreciated. Consumers can choose to transfer funds from investment instruments to cash by saving or dis-saving, which we denote s_t which is > 0 for sav-

[7]The distinction between "available cash" and "invested assets" for the sake of this model has to do with the degree to which it can be readily used as liquidity for consumption.

[8]$\zeta_t(\cdot, \cdot)$ is a function with two arguments because it also depends on the total expenditure using the debt instrument in the given period, which we will define below.

ing and < 0 for dis-saving. Labor income enters exogenously according to the function $w_t(\boldsymbol{\eta}_t)$ which depends on a vector of consumer characteristics capturing both employment status and labor productivity among others. Assume the "planner" knows $w_t(\boldsymbol{\eta}_t)$ at the time of making a period t plan, but does not know the future sequence $\{w_\tau(\boldsymbol{\eta}_\tau)\}_{\tau>t}$.

Each period, consumers make two kinds of choices: 1) how much to spend out of their various liquidity sources, 2) how much to save and pay down debt. The choice process takes a two-step approach. First, consumers decide how much of their expenditure they want to come from the different liquidity sources for each commodity each period. Second, they go out and consume and attempt to stick to their pre-defined budgets, though fail to do so with some degree of error. Let $\theta_{z,t}$ represent the share of available cash used for consumption expenditure in period t, and let $\theta_{b,t}$ be the analog share of available credit. Assume that consumers cannot spend directly out of investment assets, but must instead first transfer these assets to the cash account using s_t. Let p_t be the price of consumption. Each period, total consumption expenditure must satisfy

$$p_t c_t = \theta_{z,t} z_t + \theta_{b,t}(m_t - Q_t b_t) \tag{3}$$

Let $\epsilon_{z,t}$, $\epsilon_{b,t}$, and $\epsilon_{a,t}$ capture consumers' ex-post inattention to their ex-ante desired budgets. Balances evolve according to the laws of motion:

$$z_{t+1} = \epsilon_{z,t}\Big(w_t(\boldsymbol{\eta}_t) + (1-\theta_{z,t})z_t - e_t - s_t\Big) \tag{4}$$

$$b_{t+1} = \epsilon_{b,t}\Big((1-\theta_{b,t})(m_t - Q_t b_t) - e_t - \zeta_t\big(e_t, \theta_{b,t}(m_t - Q_t b_t)\big)\Big) \tag{5}$$

$$a_{t+1} = \epsilon_{a,t}\Big(R_t a_t + s_t\Big) \tag{6}$$

For convenience, collect the shocks into a column vector $\boldsymbol{\epsilon}_t = (\epsilon_{z,t}, \epsilon_{b,t}, \epsilon_{a,t})^\top$. Further discussion of what these shocks represent, specifically with regard to how they fit into the consumer's mental accounting decision process, is featured below.

Each period, prior to realizing the inattention shocks, the planner forms budgets to solve a recursive dynamic consumption/savings problem with nested expectations:

$$\begin{aligned} V_t(z_t, b_t, a_t, m_t) = &\max_{\substack{\theta_{z,t},\theta_{b,t} \\ e_t, s_t}} \mathbb{E}_{\boldsymbol{\epsilon}_t}\Big\{u(c_t) + \beta \mathbb{E}_t V_{t+1}(z_{t+1}, b_{t+1}, a_{t+1}, m_{t+1})\Big\} \\ \text{subject to} \quad &(3) \;\; \text{thru} \;\; (6) \end{aligned} \tag{7}$$

Note that z_t, b_t, and a_t are endogenous state variables, while m_t is exogenous, evolving according to some known stochastic process. β describes the consumer's degree of time preference. The outer expectation is taken over $\boldsymbol{\epsilon}_t$, while the expectation over future values is taken over the evolution of m_{t+1}, the evolution of η_{t+1} affecting income, and future prices p_{t+1}, R_{t+1}, and Q_{t+1}. Let $\lambda_{j,t}$, $j \in \{z, b, a\}$ be the Lagrange multiplier for the respective balance evolution constraint. The planner's ex-ante budgets must satisfy

$$\mathbb{E}_{\boldsymbol{\epsilon}_t}\left\{u'(c_t)\frac{z_t}{p_t} - \lambda_{z,t}\epsilon_{z,t}z_t\right\} = 0 \tag{8}$$

$$\mathbb{E}_{\boldsymbol{\epsilon}_t}\left\{u'(c_t)\frac{m_t - Q_t b_t}{p_t} - \lambda_{b,t}\epsilon_{b,t}(m_t - Q_t b_t)\right\} = 0 \tag{9}$$

$$\mathbb{E}_{\boldsymbol{\epsilon}_t}\left\{ - \lambda_{z,t}\epsilon_{z,t} - \lambda_{b,t}\epsilon_{b,t}\right\} = 0 \tag{10}$$

$$\mathbb{E}_{\boldsymbol{\epsilon}_t}\left\{ - \lambda_{z,t}\epsilon_{z,t} + \lambda_{a,t}\epsilon_{a,t}\right\} = 0 \tag{11}$$

For the ex-post consumption/savings problem, simply drop the expectation over $\boldsymbol{\epsilon}_t$, so that the "doer" — the side of the consumer's personality that engages in the consumption expenditure — has realized $\boldsymbol{\epsilon}_t$ chooses $\theta_{z,t}$, $\theta_{b,t}$, e_t, and s_t so as to satisfy

$$u'(c_t) = p_t\lambda_{z,t}\epsilon_{z,t} \tag{12}$$

$$u'(c_t) = p_t\lambda_{b,t}\epsilon_{b,t} \tag{13}$$

$$\lambda_{z,t}\epsilon_{z,t} = -\lambda_{b,t}\epsilon_{b,t} \tag{14}$$

$$\lambda_{z,t}\epsilon_{z,t} = \lambda_{a,t}\epsilon_{a,t} \tag{15}$$

At this point, it is useful to define an appropriate equilibrium concept that characterizes the consumer's decisions under the above mental accounting formulation.

3.2 Individual Fungibility

As was mentioned, the definition of fungibility used in this model is slightly different than that in Thaler (1985). This is formally laid out in Definition 1.

Definition 1: A consumer who exhibits "rational fungibility" has marginal expectations for the inattention shocks such that:

$$\mathbb{E}_t(\epsilon_{z,t}) = \mathbb{E}_t(\epsilon_{b,t}) = \mathbb{E}_t(\epsilon_{a,t}) = 1 \tag{16}$$

Of course, not all consumers may exhibit ex-ante rational fungibility. Some consumers may be ex-ante strongly averse to carrying unsecured debt, while others may care less. Loss aversion, as originally described in Kahneman and Tversky (1979), is a common feature built into models of mental accounting, going back to Thaler (1985). Loss aversion would manifest itself in this model if the absolute value of the dis-utility of debt is greater than the marginal utility of cash or savings. Lemma 1 highlights an important feature of the model that will be used to formally define loss aversion.

Lemma 1: Assume all decisions reside on the equilibrium path. Let dx denote a total differential over some variable x. Then the following must hold:

$$\frac{\mathrm{d}V_t^*}{\mathrm{d}z_{t+1}^*} = \lambda_{z,t}^* \epsilon_{z,t} \tag{17}$$

$$\frac{\mathrm{d}V_t^*}{\mathrm{d}b_{t+1}^*} = -\lambda_{b,t}^* \epsilon_{b,t} \tag{18}$$

$$\frac{\mathrm{d}V_t^*}{\mathrm{d}a_{t+1}^*} = \lambda_{a,t}^* \epsilon_{a,t} \tag{19}$$

From Lemma 1 it becomes apparent that the cognitive frictions are just deviations from the rationally fungible outcome, as explained in Corollary 1. Ex-ante, all consumers expect the following to hold:

$$\mathbb{E}_{\epsilon_t} \frac{\mathrm{d}V_t^*}{\mathrm{d}z_{t+1}^*} = -\mathbb{E}_{\epsilon_t} \frac{\mathrm{d}V_t^*}{\mathrm{d}b_{t+1}^*} \tag{20}$$

Thus, the ex-ante expected marginal utility of cash is equal to the negative ex-ante expected marginal dis-utility of additional debt. Consumers are making consumption decisions by adjusting balances on the margin as opposed to simply saving when cash is positive, and borrowing when no cash remains. This represents a departure from liquidity-constraint models, while still rationalizing simultaneous savings and debt utilization.

How this condition holds depends on the underlying distribution of $\boldsymbol{\epsilon}_t$. A rationally fungible consumer expects his cash balances and debt balances to move one-for-one in opposite directions, since he expects $\epsilon_{z,t} = \epsilon_{b,t} = 1$. This is summarized in Corollary 1,

Corollary 1: A consumer who exhibits ex-ante rational fungibility expects

$$\mathbb{E}_{\epsilon_t} \lambda_{z,t}^* = -\mathbb{E}_{\epsilon_t} \lambda_{b,t}^* \tag{21}$$

Proof. When $\epsilon_{z,t} = 1$, $\frac{\mathrm{d}V_t^*}{\mathrm{d}z_{t+1}^*} = \lambda_{z,t}^*$ from (17). The same holds for $\epsilon_{b,t}$, giving us the result. □

Other behavioral profiles can be defined and tested against the baseline rationality condition. Thus, $\lambda_{z,t}^*$ is the equilibrium rationally fungible marginal value of additional cash with $-\lambda_{b,t}^*$ the debt analog. With these objects in hand, a testable hypothesis regarding the model's accommodation of cash/debt utilization can be stated:

$$H_0 : \text{Ex-ante, consumers are rationally fungible.} \tag{22}$$

To test this hypothesis, a test statistic can be constructed by exploiting Lemma 1 and Corollary 1. This is defined in Proposition 1.

Proposition 1: If a consumer is ex-ante rationally fungible then

$$\mathbb{E}_{\epsilon_t} \frac{\mathrm{d}b_{t+1}^*}{\mathrm{d}z_{t+1}^*} = -\mathbb{E}_{\epsilon_t} \frac{\epsilon_{b,t}^*}{\epsilon_{z,t}^*} = -1 \tag{23}$$

Proof. Under the ex-post analog of Corollary 1, $\frac{\mathrm{d}V_t^*}{\mathrm{d}z_{t+1}^*} = \lambda_{z,t}^*$, and the same for debt. Thus we can write

$$\lambda_{z,t}^* = -\lambda_{b,t}^* \tag{24}$$

since $\epsilon_{z,t} = \epsilon_{b,t} = 1$. This implies that

$$\frac{\lambda_{z,t}^*}{\lambda_{b,t}^*} = -1 \tag{25}$$

Now substitute the total derivatives back in and cancel $\mathrm{d}V_t^*$ terms to get

$$\frac{\mathrm{d}b_{t+1}^*}{\mathrm{d}z_{t+1}^*} = -1 \tag{26}$$

Add back the expectation operator, and the result is had. □

Of course, ex-post rational fungibility is a measure zero outcome. The null hypothesis that a consumer exhibits rational fungibility can be tested using (23). Balance evolutions are observed. Let Δ be a suitable finite differences operator and let $\widehat{\mathbb{E}}$ denote a finite sample mean operator. Then testing H_0 amounts to testing:

$$H_0 : \widehat{\mathbb{E}}_{\epsilon_t} \frac{\Delta b_{t+1}^*}{\Delta z_{t+1}^*} + 1 = 0 \tag{27}$$

For now, assume all consumers are willing to hold some level of debt, so that debt holdings are a strictly monotonic function of cash holdings.[9] Then two possible alternative hypotheses can be associated with H_0. Consumers can either be debt averse, or debt loving. Debt aversion is the model analog of classic mental accounting and prospect theory debt aversion, while debt loving behavior is its antithesis. Definitions 2 and 3 formalize these alternative hypotheses.

Definition 2: A consumer is said to be "debt averse" if either

i. Debt decreases more than cash increases.

ii. Cash decreases more than debt increases.

iii. Debt decreases and cash decreases simultaneously, but debt decreases more than cash decreases.

iv. Cash increases and debt increases simultaneously, but cash increases more than debt increases.

Definition 3: A consumer is said to be "debt loving" if either

[9] It should probably also be noted that all consumers need to have some kind of checking or cash equivalent account which they use to make payments on balances due.

i. Cash increases more than debt decreases.

ii. Debt increases more than cash decreases.

iii. Debt decreases and cash decreases simultaneously, but cash decreases more than debt decreases.

iv. Cash increases and debt increases simultaneously, but debt increases more than cash increases.

In the event that a consumer refuses to hold any debt, then his behavior falls under Definition 4.

Definition 4: A consumer is said to be "absolutely debt averse" if $b_t^* = 0, \forall t$.

Propositions 2 and 3 formally summarize the behavioral profiles associated with definitions 2.*i*. thru 2.*iv*. and 3.*i*. thru 3.*iv*..

Proposition 2: For a consumer who is debt averse $b_{t+1}^* \geq 0$ and one and only one of the following must hold on average:

$$\widehat{\mathbb{E}}_{\epsilon_t} \frac{\Delta b_{t+1}^*}{\Delta z_{t+1}^*} + 1 < 0 \tag{28}$$

$$\text{or} \qquad \Delta b_{t+1}^*, \Delta z_{t+1}^* < 0 \quad \text{and} \quad |\Delta b_{t+1}^*| > |\Delta z_{t+1}^*| \tag{29}$$

$$\text{or} \qquad \Delta b_{t+1}^*, \Delta z_{t+1}^* > 0 \quad \text{and} \quad |\Delta b_{t+1}^*| < |\Delta z_{t+1}^*| \tag{30}$$

Proposition 3: If a consumer is debt loving then $b_{t+1}^* \geq 0$ and:

$$\widehat{\mathbb{E}}_{\epsilon_t} \frac{\Delta b_{t+1}^*}{\Delta z_{t+1}^*} + 1 > 0 \tag{31}$$

Proposition 2 describes all three possible conditions associated with a debt averse consumer. Meanwhile, Proposition 3 provides a necessary condition for a consumer to be debt loving, though this condition is not sufficient. Note that when either $\Delta b_{t+1}^*, \Delta z_{t+1}^* < 0$ or $\Delta b_{t+1}^*, \Delta z_{t+1}^* > 0$ condition (31) holds, but the consumer can still be debt averse, as described in Proposition 2. Meanwhile, the set of consumers who are absolutely debt averse does not intersect with the set of consumers who are debt averse. This is fairly straightforward to see as illustrated by Proposition 4.

Proposition 4: If a consumer is absolutely debt averse then on average $\mathbb{E}_{\epsilon_t} \Delta b_{t+1}^* = 0$ and

$$\widehat{\mathbb{E}}_{\epsilon_t} \frac{\Delta b_{t+1}^*}{\Delta z_{t+1}^*} + 1 = 1 \tag{32}$$

3.3 Aggregate Fungibility

If a researcher possesses a panel of time series data for credit card and cash balances, then the degree to which consumers are generally debt averse or debt loving can be quantified and (27) can be tested. In fact, we can test both whether a particular individual behaves on average according to (27) and whether on average all individuals behave according to (27). Let subscript i index individual balances and let $\widehat{\mathbb{E}}_{\mathcal{I}}$ denote the finite sample mean over the set of all individual consumers for whom researchers have observations in period t, $\mathcal{I}_t$. Let $\xi_{i,t}$ denote a particular observation of the fungibility condition so that

$$\xi_{i,t}(\epsilon_{i,t}) = \frac{\Delta b^*_{i,t+1}}{\Delta z^*_{i,t+1}} + 1 \tag{33}$$

Using (33), we can define aggregate conditions on the distributions of consumers with different behavioral profiles.

Consider all consumers who are either debt averse or absolutley debt averse and all consumers who are debt loving. In Definition 5, we describe conditions on the aggregate distribution of consumers.

Definition 5: Let $\mathcal{I}_{\text{Averse}} \subseteq \mathcal{I}$ denote the set of consumers who are either debt averse or absolutely debt averse, and let $\mathcal{I}_{\text{Loving}} = \mathcal{I}^c_{\text{Averse}}$ be the compliment of said set. We say on aggregate that:

i. On average, consumers are debt averse if

$$\left| \widehat{\mathbb{E}}_{\mathcal{I}} \left\{ \widehat{\mathbb{E}}_{\epsilon_{i,t}} \xi_{i,t}(\epsilon_{i,t}) \;\middle|\; i \in \mathcal{I}_{\text{Averse}} \right\} \right| > \left| \widehat{\mathbb{E}}_{\mathcal{I}} \left\{ \widehat{\mathbb{E}}_{\epsilon_{i,t}} \xi_{i,t}(\epsilon_{i,t}) \;\middle|\; i \in \mathcal{I}_{\text{Loving}} \right\} \right| \tag{34}$$

ii. For debt loving, flip the sign in (34).

Consumers are on average debt averse if the absolute value of the expected error for the whole set of debt averse consumers is greater than the absolute value of the expected error for the whole set of debt loving consumers. Definition 4 can be used as a rule of thumb to better understand aggregate behavior. Integrating over the entire distribution of consumers, note that a few extremely debt averse consumers will push the absolute value of the mean of the conditional distribution rightward. Clearly we can also count the number of consumers who fall under each condition, thus assessing different moments of the underlying distribution. This information can be useful when running simulations to understand how the joint distribution of income, wealth, and debt will evolve in response to both aggregate and idiosyncratic shocks.

4 Identification & Estimation

The distribution of fungibility both for each consumer and across all consumers is sought. Preliminary estimation of a similar test of fungibility derived with different timing assumptions as to how interest accrues on debt balances can be found in Montgomery, Olivola, and Pretnar (2018). To estimate the distribution over consumers of $\mathbb{E}_{\epsilon_{i,t}}\xi_{i,t}(\epsilon_{i,t})$, researchers require balance evolution information about both debt-bearing consumer credit accounts and a checking account or its cash equivalent. Given this information and nothing more, (27) can be tested consumer by consumer and aggregate fungibility can be assessed using (34) or its debt loving reversed form. However, such estimates will be biased, failing to account for how consumers plan for expected income and price shocks or recurring purchases, like utility bills and insurance premiums. This section discusses both the identification issues associated with estimation and proposes ways to ameliorate them, controlling for various biases.

4.1 Identification

Foremost, each observation of $\xi_{i,t}(\epsilon_{i,t})$ could be influenced by ex-ante anticipated changes to income and prices. For example, if consumers get a low income shock one period, and anticipate it, then they may ex-ante intentionally plan to spend more on their credit card beyond what fits the condition. The fungibility statistic is thus income dependent. Indeed, one could say that income changes, or deviations from median/mean income are correlated with the fungibility statistic but exogenous to the consumer's underlying ability to ex-post stick with planned consumption budgets. Researchers can control for income changes directly if income data is possessed. Anticipated price change affect estimates of the distribution for $\xi_{i,t}(\epsilon_{i,t})$ similarly. One can imagine how a consumer may plan his upcoming expenditure profile expecting prices to be at a certain level, but goes to the store and is surprised to find that the price of a particular high-dollar purchase has gone up substantially. Needing to make this purchase, he decides to put it on his credit card, whereas he he may have otherwise intended to pay cash. His decision thus looks exactly like a cognitive lapse in his ex-ante expenditure planning, but it is utterly intentional. Researchers can control for price variation around a long-run trend in several categories, using several broad indices.

Additionally, there remains to deal with both recurring purchases and unanticipated emergency purchases. Recurring purchases, like for example, bill payments that are structured in fixed, monthly amounts should be fully baked into consumers' anticipated balance changes, since they know ahead of time that such payments will be automatically deducted. Thus, it is assumed that consumers have no cognitive dissonance with regards to their recurring purchases, so that any deviation from the rational fungibility condition is itself independent of the presence of recurring purchases. Unanticipated emergency expenditures however, like medical expenses resulting from an accident or some forms of automobile maintenance, themselves represent shocks to the consumer's consumption path, initiating unanticipated deviations from a planned budget.Thus, an unanticipated emergency expenditure is in itself a fungibility shock. If over the long run, all emergency expenditures are placed on a credit card and balances between credit cards and checking accounts are not adjusted accordingly, then the behavioral profile of the consumer will on average deviate from rational fungibility. After all, he is failing to formulate rationally fungible plans, otherwise his behavior would not deviate so strongly from the intratemporal balance smoothing condition we derive. This does not, however, preclude identification of his underlying behavioral tendency. Rather, in fact, it enhances it: if he were truly rationally fungible, then he would be able to adjust his balances over the long run, despite the presence of emergency expenditures, in order to satisfy rational fungibility.

4.2 Estimation Strategy

The atomic unit of observation is individual i's balance evolution from period t to $t+1$. Only debt averse of debt loving consumers must be considered: absolutely debt averse consumers have all zero observations and so, for them, the effects of anticipated income changes and unanticipated price changes cannot be identified. The dataset used to estimate a similar test of fungibility in Montgomery, Olivola, and Pretnar (2016) is unbalanced, which can potentially be problematic if a consumer-level fixed effects model is intended. For example, if a consumer decides to close a credit card or stop using it altogether, rather than just falling from the observation randomly, then fixed effects estimates will be biased, and the direction of such bias is difficult to ascertain. Why might a consumer stop using a credit card? Perhaps, he is regulating his consumption expenditure and quitting credit cards cold turkey is, he feels, the best way to accomplish this. He could also simply lose the card, pay off the balance, and never bother to replace it, instead using one of the other credit card accounts he has open. Such an account is effectively closed even though from the bank's perspective, it is simply a credit card with persistent zero balance that is never used. With unbalanced data, researchers thus face a choice as to whether to estimate separate regression models for each agent, controlling for income and price variation, or dealing with the potential endogeneity bias wrought by estimating a standard fixed effects model. The necessary control variables for the regressions posed here are functions of only income and prices. If a researcher were to run a standard fixed effects regression on unbalanced panel data, he could assume that income variation and price variation effect agent-level fungibility outcomes uniformly. But such a model structure would seem to be counterintuitive to the underlying goal of the quantitative exercise: to estimate both individual distributions of fungibility and the aggregate distribution.

Given these issues, it is best to sacrifice degrees of freedom robustness and estimate separate regression models for each agent. Let $\overline{w}_i$ be mean income for agent i. Let p_t be am aggregate price realization, like say from the Consumer Price Index (CPI) or Personal Consumption Expenditure (PCE) index. Let $\overline{p}_{N,t}$ be a simple moving average price realization over N backwards lags. Let σ_{w_i} be the standard deviation of income, and let $\sigma_{p_{N,t}}$ be the standard deviation of prices over the N periods of backwards lags. We seek to understand how deviations from expected income and price effect fungibility outcomes. Denote agent i standardized income for period t as $\widetilde{w}_{i,t} = \frac{w_{i,t}-\overline{w}_i}{\sigma_{w_i}}$. Denote period t standardized prices over an N-lag moving average as $\widetilde{p}_t = \frac{p_t-\overline{p}_{N,t}}{\sigma_{p_{N,t}}}$. The problem formerly posed then is to estimate for each i:

$$\xi_{i,t}(\epsilon_{i,t}) = \beta_{i,0} + \beta_{i,1}\widetilde{w}_{i,t} + \beta_{i,2}\widetilde{p}_t + \nu_{i,t} \tag{35}$$

Under our identification assumptions $\mathbb{E}[\nu_{i,t} \mid \widetilde{w}_{i,t}, \widetilde{p}_t] = 0$.

Controlling for income changes in this manner assumes that agents know exactly how much they will earn. This is a reasonable assumption: labor income, for example, is deposited in a checking account or cut through a check two to four weeks, on average, after the income is earned. It is reasonable to assume agents know the value of their upcoming paychecks. Having thus controlled for anticipated income changes and price shocks, the regression analog of the hypothesis in (27) is

$$H_0 : \beta_{i,0} = 0 \tag{36}$$

This hypothesis can be tested with a simple t-test on the regression coefficient, and a consumer's behavioral profile with respect to consumer debt can then be subsequently verified. From there, the sample of debt averse or absolutely debt averse consumers can be defined and the aggregate nature of the sample, as described in Definition 5, can be verified.

5 Moving Forward

Moving forward the MPC regressions and the structural model will be estimated using the data described in Montgomery, Olivola, and Pretnar (2018). Once the distribution of structural parameters is found, counterfactual shocks can be simulated to assess how the joint distribution of income, wealth, debt, and consumption evolves. The counterfactual simulation strategy will involve simulating both net-zero idiosyncratic income and credit shocks and net-positive and net-negative aggregate income and credit shocks. All results, including counterfactuals, will be linked to the household's zip code. Analysis of household behavior can then be assessed against neighborhood demographic and socioeconomic characteristics to understand if the broader economic environment may be acting upon consumer decisions.

References

Browning, Martin, and Annamaria Lusardi. 1996. "Household Saving: Micro Theories and Micro Facts". *Journal of Economic Literature* 34:1797–1855. (Cit. on p. 2).

Gross, David, and Nicholas Souleles. 2002. "Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data". *The Quarterly Journal of Economics* 117 (1). (Cit. on p. 2).

Kahneman, Daniel, and Amos Tversky. 1979. *Econometrica* 47 (2): 263–292. (Cit. on p. 6).

Laibson, David, Andrew Repetto, and Jeremy Tobacman. 2000. "A Debt Puzzle". *NBER Working Paper #7879*. (Cit. on p. 4).

Montgomery, Alan, Christopher Olivola, and Nick Pretnar. 2016. "A Structural Model of Mental Accounting". *Working Paper: www.npretnar.com*. (Cit. on p. 12).

— . 2018. "Durables, Non-Durables, and a Structural Test of Fungibility". *Working Paper, www.npretnar.com*. (Cit. on pp. 2, 11, 13).

Telyukova, Irina. 2013. "Household Need for Liquidity and the Credit Card Debt Puzzle". *The Review of Economic Studies* 80. (Cit. on pp. 1, 4).

Telyukova, Irina, and Randall Wright. 2008. "A Model of Money and Credit, with Application to the Credit Card Debt Puzzle". *The Review of Economic Studies* 75. (Cit. on pp. 1, 4).

Thaler, Richard. 1985. "Mental Accounting and Consumer Choice". *Marketing Science* 4 (3): 199–214. (Cit. on pp. 3, 4, 6).