

Why Has Consumer Spending Remained So Resilient? Evidence from Credit Card Data^{*}

Rees Hagler[†]

Dhiren Patki[‡]

August 2025

Supplementary Appendix

^{*}Any opinions and conclusions expressed herein are those of the authors and do not indicate concurrence by the Federal Reserve Bank of Boston, the principals of the Board of Governors, or the Federal Reserve System. All errors are our own.

[†]Federal Reserve Bank of Boston. Email: Rees.Hagler@bos.frb.org

[‡]Federal Reserve Bank of Boston. Email: Dhiren.Patki@bos.frb.org

A Income Imputation

In this Appendix, we describe our income-imputation methodology for the FR Y-14M data, which, we note, is experimental and is being actively improved. We provide an overview of the estimation procedure here and then show validation statistics.

Account-holder household income is reported in two conceptually distinct cases in our data. The first is at account origination when we observe household income for account holders through their credit applications.¹ After origination, we observe about 40 percent of account holders providing updates of their household income through 2020, after which post-origination updates are no longer available in Y-14M data. In our data, we observe self-reported income—either at origination or as a post-origination update—for 22 percent of card-year observations. We use this subset of observations to develop our imputation procedure as follows:

1. We convert nominal income values into 2022 dollars using the Personal Consumption Expenditures (PCE) price index. We then divide this observed card-year income distribution into quintiles, which we also refer to as income groups.
2. We convert all other nominal dollar values, such as spending, credit limits, and revolving debt balances, into 2022 dollars using the PCE price index.
3. We split our data into a training sample representing 90 percent of our observations and a test sample representing 10 percent of our observations.
4. Using the training sample, we employ a random forest model to predict the updated income quintile conditional on the following variables:
 - (a) Income quintile at origination
 - (b) Log average real monthly spending in the year of the income update
 - (c) Account age in months
 - (d) State of residence of the account holder
 - (e) Account holder type (that is, joint or single)
 - (f) Real credit limit quintile at origination
 - (g) Real credit limit quintile in the year of the income update
 - (h) Indicator for 2020 (to account for COVID-19 pandemic effects)
 - (i) Indicator for whether the account carries a positive revolving debt balance
 - (j) Log real revolving debt balance quintile in the year of the income update
 - (k) Indicator for whether the account holder has multiple banking relationships (that is, the credit card account holder may also have a checking account, mortgage, etc. with the issuing bank).

¹Income at origination is missing for about 25 percent of the credit cards. We assign these cards to the lowest income group when implementing our imputation procedure.

The random forest model captures nonlinearities and interactions between the variables, which aid in prediction. In Table A1, we evaluate the performance of the model on the test sample by comparing the predicted income quintiles (reported in the columns) to the actual income quintiles (reported in the rows). About 50 percent of the observations are correctly classified in each quintile, with slightly higher shares at the very bottom and very top. Misclassifications are heavily concentrated in adjacent quintiles, with about 70 to 75 percent of the sample being assigned a level no more than one quintile above or below its true level.

Imperfect predictions have the effect of masking disparities between the different income groups. For instance, Table A1 shows that while about 57 percent of accounts in the first quintile are correctly classified, it is also the case that about 5 percent of accounts from the fifth quintile are incorrectly assigned to the lowest income group. As a result, outcomes for individuals in the first quintile are pulled toward those from the fifth quintile. In general, these classification errors imply that outcomes for each predicted income group end up looking more similar to the population mean than the true (but unobserved) outcomes. The between-income-group differences that we show in our brief are therefore conservative estimates of the true gaps in outcomes.


Table A1: Validation of Income-Quintile Prediction

Actual Income Quintile ↓	Predicted Income Quintile →				
	1	2	3	4	5
1	0.567	0.195	0.118	0.068	0.052
2	0.161	0.496	0.195	0.095	0.054
3	0.083	0.168	0.491	0.173	0.085
4	0.060	0.087	0.181	0.485	0.188
5	0.065	0.077	0.121	0.212	0.525

Notes: This table shows the out-of-sample performance of the random forest prediction model in the test sample. Diagonal entries show the share of card-years within a given income quintile that are correctly classified. Off-diagonal entries show the respective share that are misclassified.

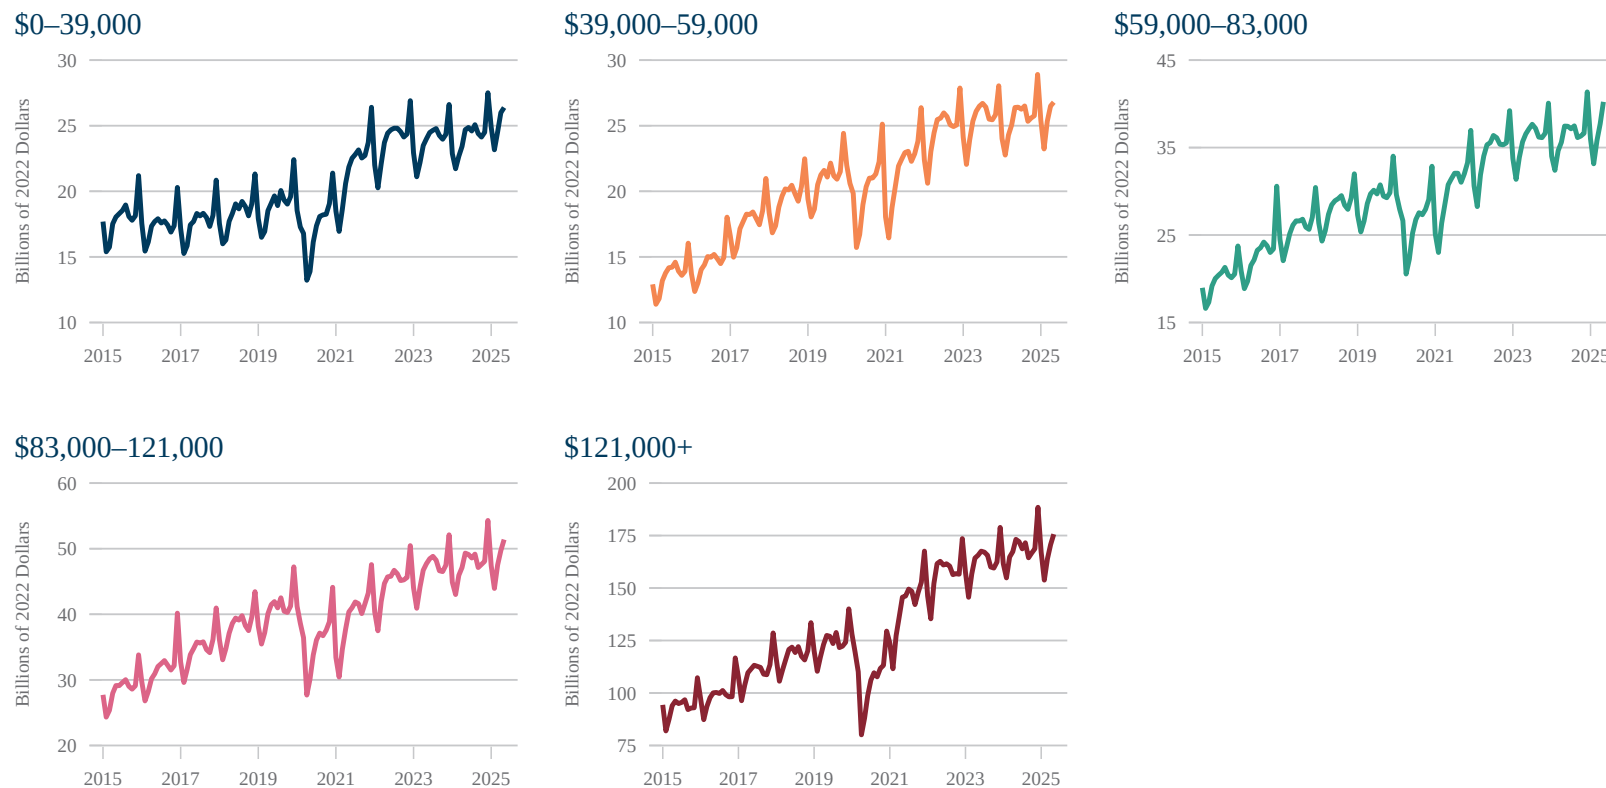
B Additional Figures

Figure B1: Aggregate Spending in Y-14M Compared with the Census Monthly Retail Trade Survey (MRTS)



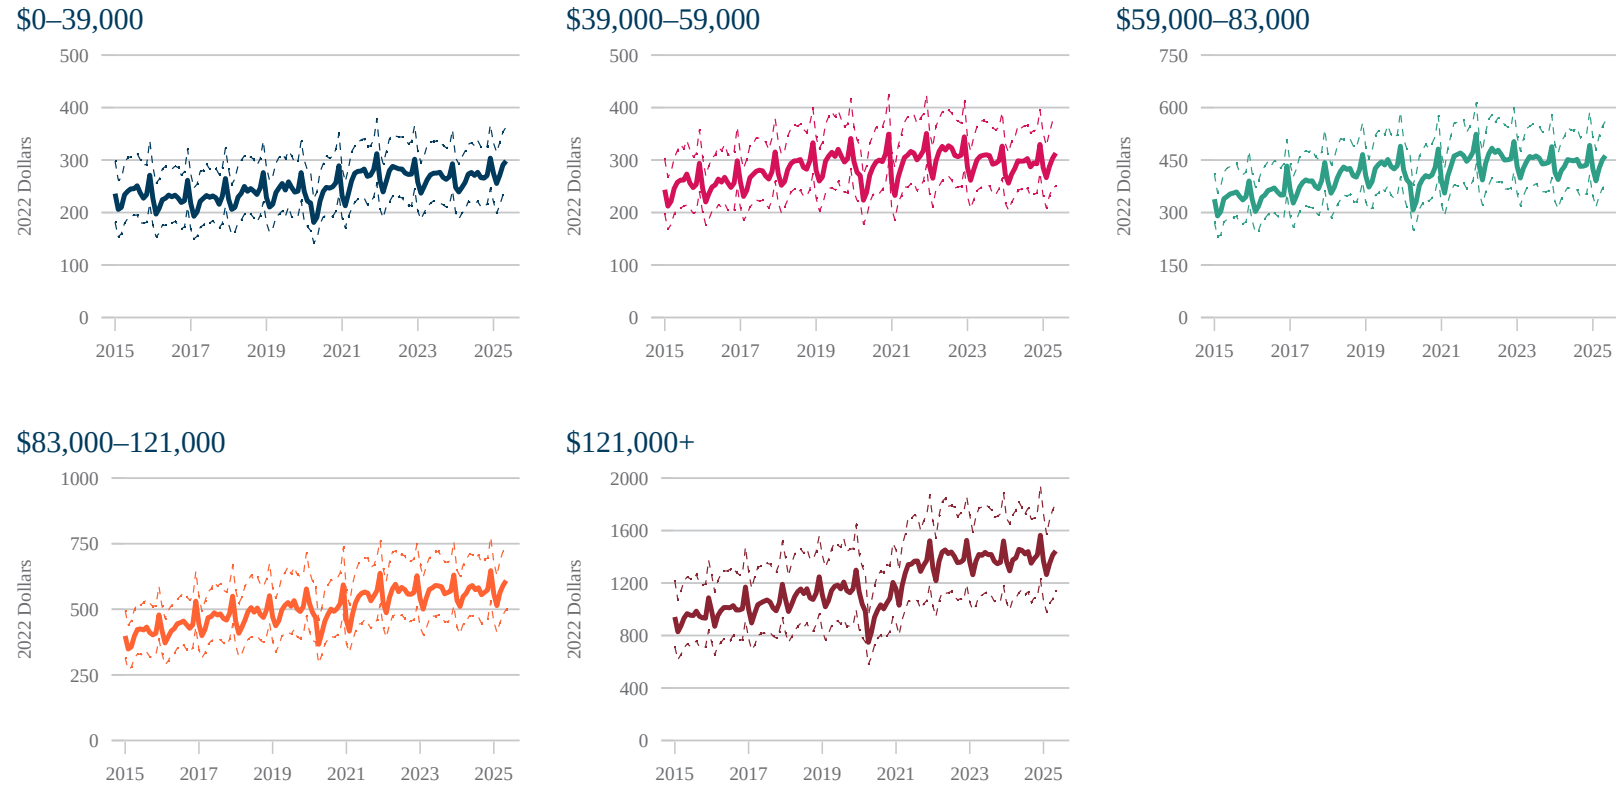
Note: Values adjusted to 2022 Dollars using the PCE price index. MRTS through June 2025. Y-14M through May 2025.
Source: Y-14M, Census MRTS / Haver Analytics.

Figure B2: Aggregate Spending in Y-14M, by Income Group




Note(s) and Source(s): January 2015–May 2025. Y14M and author's calculations.

Figure B3: Spending per Card by Income Group and County



Note(s) and Source(s): January 2015–May 2025. Y14M and author's calculations.

Figure B4: Aggregate Credit Card Debt by Income Group



Note(s) and Source(s): January 2015–April 2025. Y14M and author's calculations.