Rainy Day Credit?
Unsecured Credit and Local Employment Shocks*

Benjamin J. Keys
The Wharton School
University of Pennsylvania

Jeremy Tobacman
University of Delaware

Jialan Wang
University of Illinois
at Urbana-Champaign

March 2018

Abstract

Aggregate credit card balances fell by 15% during the Great Recession and have not yet surpassed their prior peak. This paper investigates the causes of the national cycle by exploiting geographic variation in the intensity and timing of the recession and recovery. Specifically, we instrument for local changes in employment using a Bartik (1991) style methodology, based on pre-period county-level industrial composition interacted with nationwide industry trends. Using a high-quality dataset covering a large fraction of U.S. credit card accounts, we find that both purchase volumes and payment volumes increase in response to plausibly exogenous positive employment shocks, while credit limits and balances decrease. We attempt to reconcile these findings with theories of consumption decision-making and with nationwide aggregates. Countercyclical demand for credit card balances implies that procyclical credit supply responses are larger than previously estimated, and amplify rather than mitigate consumption volatility driven by the business cycle.

*We would like to thank audiences at the CFPB, Delaware, Drexel, the FDIC Consumer Research Symposium, Wharton, the Philadelphia Fed, and the ASSA meetings for helpful feedback and Min Kim for outstanding research assistance. The views expressed are those of the authors and do not necessarily reflect the opinions of the Consumer Financial Protection Bureau, its staff, or the US Government. Any remaining errors are our own. Contact: benkeys@wharton.upenn.edu, tobacman@udel.edu, jialanw@illinois.edu.

I Introduction

How does the credit card market respond to local employment shocks? Standard consumption models imply that credit demand will rise in response to adverse transitory income shocks. However, the effect of adverse shocks on equilibrium credit utilization is ambiguous, because adverse shocks increase default risk and lenders may have limited capacity to distinguish between transitory and persistent shocks to borrower creditworthiness.

In this paper, we seek to investigate and quantify how much self-insurance via credit card borrowing was demanded and supplied during the period between 2008 and 2014. In the Great Recession, as unemployment rates rose sharply, revolving credit outstanding in the United States fell 15% from its June 2008 (nominal, seasonally-adjusted) peak (Federal Reserve G19 Series). In the aggregate, this implies that the supply response by lenders outweighed the heightened credit demand as employment fell. The aggregate patterns, however, belie considerable heterogeneity at the county level. We exploit variation in county-level employment shocks, using a Bartik-style industry composition shift-share instrument, in order to produce new estimates of the elasticity of equilibrium credit card accounts, limits, interest rates, and balances with respect to income and employment.

The estimation and interpretation of causal effects using Bartik-style instruments to isolate shocks to labor demand has considerable history by now, for a variety of purposes.[1] To our knowledge, this is the first work to use this sort of variation to study the effects of income shocks on consumer credit utilization.

Our primary goal is to understand whether private sources of consumer credit facilitate self-insurance. The cyclicality of unsecured credit may either amplify or attenuate recessions, depending on the responsiveness of credit supply and credit demand. While an extensive literature has explored government support for the financial and household sectors during the Great Recession, our paper seeks to examine the private credit market's support for households during this contractionary period.

[1] An incomplete list of papers that have used this approach, or one similar, includes Murphy and Topel (1987), Blanchard and Katz (1992), Bound and Holzer (2000), Gould, Weinberg and Mustard (2002), Autor and Duggan (2003), Aizer (2010), Greenstone, Mas and Nguyen (2015), Hoynes, Miller and Schaller (2012), Diamond (2016), Chodorow-Reich, Feiveson, Liscow and Woolston (2012), Maestas, Mullen and Powell (2016), and Notowidigdo (2013).

We take advantage of extraordinarily high-resolution data on credit card accounts from the CFPB Credit Card Database (CCDB), which has coverage of 90% of all accounts nationwide. We construct county-quarter measures of credit card accounts, credit limits, payments, purchases, revolving balances, and costs (APR). These rich measures of credit card usage provide an unprecedented insight into the inflows and outflows of the unsecured credit market during the Great Recession.

Using "long-difference" specifications that rely solely on cross-sectional variation across counties in employment shocks, we find that counties with larger adverse unemployment shocks increased their number of accounts, decreased their payments, and increased their overall outstanding revolving balances.

Notably, the aggregate patterns during our period of analysis go in the opposite direction from the elasticities estimated off of well-identified cross-sectional variation. Large declines in employment nationwide during the Great Recession was associated with large decreases in balances, a pattern that holds in many markets.

How can these discrepancies best be reconciled? The most straightforward explanation is that there were simultaneous shocks to the credit supply side of the market. The period we study, the Great Recession and its aftermath, was characteristics by the broader financial crisis and large declines in bank lending. In short, we find that countercyclical credit demand was massively dominated by procyclical credit supply.

The impact of tightening credit when demand was especially high had huge impacts on consumption during the Great Recession. While other work has focused on the "debt overhang" of secured debt such as mortgages, our results suggest that latent demand for unsecured credit likely far outstripped the amount of credit provided during the boom (Mian, Rao and Sufi 2013). Thus, our results complement those of Agarwal, Chomsisengphet, Mahoney and Stroebel (2016), who use a different identification approach, namely discontinuities in credit card offer algorithms, to show the disconnect between credit supply and credit demand during the post-crisis recovery period.

Our estimates provide new evidence for the relationship between employment shocks and credit use. Sullivan (2008) and Keys (2017) explore the private-side responsiveness of credit markets to income shocks. More generally, our findings underscore concerns regarding unsecured credit access, which may not be available when needed most. The self-insurance aspects of credit cards are limited when issuers respond aggressively to deteriorating economic conditions, reducing credit lines and canceling accounts. Our results suggest that if credit supply amplifies business-cycle consumption volatility rather than mitigating it through consumption-smoothing, then government programs must additionally support the household sector above and beyond what would be expected in a static framework.

In the next section, we describe the hypotheses that we seek to test. Section III introduces our primary data sources and provides descriptive statistics. Our identification strategy is presented in Section IV, and results are discussed in Section V. Section VI concludes.

II Hypotheses

Consumer credit markets respond to many factors. This paper focuses on impacts of local labor demand shocks. When a county experiences a labor demand shock, we expect to see consumer credit effects via consumer reoptimization and via issuers' reassessment of the profitability of cardholders in that county. These two effects can be complementary or offsetting. In this section, we hypothesize about the direction of net county-level effects of positive labor demand shocks on the number of accounts, credit limits, APRs, purchase volumes, payments, and balances.

One aspect of our analysis that has both advantages and limitations is that we study effects aggregated to the county level. Changes in employment in a county reflect the net effect of job gainers, job losers, existing workers who continue to work, and existing non-workers who continue not to work. The changes in employment that we measure and use aggregate across those groups In addition, changes in credit use may have spillovers. Discerning mechanisms is challenging when all this is potentially at play, but some evaluation of policy counterfactuals is best conducted inclusive of these various channels. [2]

In principle outcomes could be calculated in aggregate, per account for all accounts open in a quarter, or per account for accounts that were open in the base quarter. Some variables are easier

[2] A variety of other factors also may affect credit card offers and behavior, not least including asymmetric information, expectations about future local labor demand shocks, durable goods, consumption habits, present bias, reference dependence, and substitution across household financial products (Gorbachev and Luengo-Prado forthcoming). To the extent these can be identified, we defer discussion of them until later in the paper.

to interpret when normalized by the number of accounts, but a simple supply/demand framework connects most naturally with the county aggregates (especially under the assumption that migration is small), so we begin there. If supply increases when employment exogenously increases, Q should increase and P should decrease. If demand contracts when there's a rise in employment, then Q should decrease and P should decrease. In equilibrium, if both supply and demand shift (with S rising and D falling), then P decreases and Q is ambiguous. The closest analogues for P and Q in the CCDB are APR and balances outstanding.

We hypothesize about the direction of effects beginning with the number of accounts. When employment goes up exogenously in a county, generally we would expect the default risk of new and existing account-holders to go down, incentivizing card issuers to open more new accounts and close fewer existing accounts. For expositional simplicity, we summarize this reasoning by saying that we expect the credit supply effect on the number of accounts to be positive. The credit demand effect of a positive employment shock on the number of accounts is ambiguous. Transaction demand to make purchases could go up, while the demand for new borrowing would likely go down. Thus the net effect is ambiguous in direction, and determining its sign and magnitude is an empirical question.

We would expect the credit supply and credit demand effects on credit limits should respond similarly in direction. Turning to effects on APR, in the event of positive employment shocks, we expect default risk to fall implying a negative credit supply effect. As discussed above we expect the demand for borrowing to fall when employment rises, implying a negative credit demand effect on APRs.

Purchase volumes and payments on existing accounts may not be affected much by employment-driven changes in credit supply. When employment exogenously increases, we would expect consumer reoptimization to involve increases in both volumes and payments. The net effect on credit card balances is ambiguous as it depends on these mangitudes, the interest rate, and marginal accounts. Another way to think about the effect on balances is to view balances as the quantity of credit card credit. When supply shifts out and demand shifts in, the effect on quantity is ambiguous.

The directional ambiguity of many of these predicted effects motivates empirical investigation, to which we now turn.

III Data

Our primary dataset is the CFPB Credit Card Database (CCDB), which includes account-level data for a number of large credit card issuers in the United States. The data are collected under the CFPB's regulatory authority over the credit card market as prescribed by the Dodd-Frank Act.[3] The data used here cover February 2008 to December 2013, and the issuers in the full dataset comprise over 85% of credit card industry balances.

The dataset includes information on the near-universe of consumer and small business credit card accounts from included issuers. The variables include monthly account-level details on balances, payments, fees, interest rates, and delinquency. In addition, the CCDB includes FICO scores and individual income both at origination and based on periodic updates by issuers. Each account is linked to credit bureau data that provide a summary of the borrower's overall credit portfolio on a quarterly basis. While we cannot link separate accounts to the same consumer or household, we can observe total credit card activity for each individual (including any joint accounts with other household members) in the credit bureau variables. The CCDB does not contain data on individual purchase transactions.

We apply two primary restrictions to the full CCDB to arrive at our analysis sample. First, we restrict our sample to a balanced sample of issuers which are available from 2008-2013. We exclude any issuers that enter or exit the sample during this period due to data availability, portfolio transitions, or mergers and acquisitions. Second, to measure accounts and balances that are available for use by consumers, we only consider active accounts as flagged by the issuers. Our analysis dataset includes the full universe of active accounts from included issuers.

Starting with the microdata at the account-month level, we aggregate to the zipcode-quarter level based on the billing zipcode of each account. For stock variables such as balance, credit limit, and number of accounts, we measure the stock in each zipcode in the last month of each quarter. For flow variables such as purchase volume and payments, we take the sum over the three months

[3] The dataset also includes nine institutions that fall under the purview of the U.S. Office of the Comptroller of the Currency (OCC). For additional information on the CCDB, see CFPB (2013) and Keys and Wang (2016).

in each quarter. Our measure of APR averages over all active accounts in each cell. Finally, we merge the zipcode-quarter data to the HUD/Census zipcode-to-county crosswalk and collapse to the county-quarter level to obtain the final dataset for our analysis.[4]

Summary statistics for the CCDB are presented in Table I. Over our window of observation, the number of credit card accounts fell by 14 percent, while the average credit limit per account fell from $11,000 to $9,000. Purchases and payments increased and average balances fell on net over the period, while interest rates stayed relatively constant.

IV Identification Strategy

Our main specifications regress percentage changes in credit outcomes on percentage changes in employment, where we instrument for the employment change using Bartik-style instruments. Thus the coefficients we estimate can be interpreted as elasticities. In all the specifications we report in this paper, we winsorize the percent changes at the 1% level, and we weight by the time-0 number of credit card accounts.

We instrument for the change in a county's employment with the interaction of the pre-period industry mix of employment in that local labor market and the national change in industry employment (exclusive of the given county). These measures are constructed using the Quarterly Census of Employment and Wages (QCEW) at the county level. Identification requires that the pre-period industrial mix interacted with the national industry trends does not directly affect local credit card variables.

We focus on a 'long difference" specification that relies solely on cross-sectional variation in the changes in employment and credit card market variables. We also consider a county-quarter panel version that uses quarterly variation, but is subject to critiques regarding serially correlated errors within a geography over time.

Formally, the long difference specification reads as follows for a CCDB variable $CCDB_{ct}$ in county c between times t and t+k:

[4] For the few zipcodes that span multiple counties, we allocate the zip-level data proportionally to the counties using the residential address shares in the crosswalk.

$$\Delta log(CCDB_{c,t+k}) = \beta_0 + \beta_1 * \Delta log(Employment_{c,t+k}) + \epsilon_{c,t+k}$$

where we instrument for $\Delta log(Employment_{c,t+k})$ with a measure of predicted employment growth that equals, summing across industries i,

$$\Delta log(PredictedEmployment_{c,t+k}) = \Sigma_i \frac{Employment_{i,c,t}}{Employment_{c,t}} * \Delta log(Employment_{i,-c,t+k})$$

A complementary way to interpret the instrument is to notice that the predicted employment growth in county c between times t and t+k is a weighted average of national employment growth rates across industries during the same period of time (exclusive of county c), where the weights are the time t employment shares in county c.

Using the notation of this section, our results below generally fix t=2008q1 and plot estimates as a function of t+k. In this way we are able to estimate and compare both short- and long-run elasticities.

V Results

V.A First Stage

Our identification strategy leverages cross-sectional variation in exposure to industry-specific fluctuations in demand. We show that the industry shift-share methodology yields a strong first stage in Figure 2, which presents the relationship between the predicted change in employment on the x-axis and the actual change in employment on the y-axis. The F-test for the panel version of this regression is over 1,000.

Figure 3 shows the time-series version of quarterly sequential long-difference specifications of the first-stage. In other words, this figure plots the coefficients on 7 distinct long-difference specifications, each of which use 2008q1 as the base level of employment shares by industry. The instrument remains highly statistically significant throughout the 2009-2015 period, although the strength of the instrument falls steadily as more time elapsed.

In general, we exploit the variation in industry composition across local markets to estimate the responsiveness of credit card borrowing to employment shocks using a standard Bartik-style shift-share instrument. In what follows we focus on the instrumental variable results. The differences between the OLS and the IV are best understood as reflecting responses that are not induced by industry-specific demand shocks, in other words that may be influenced by other labor supply factors or regional trends.

One particular threat to inference that is addressed when we move from the OLS to the IV is reverse causality. Suppose that consumer credit utilization rose exogenously in a county. Plausibly, consumers would spend money locally, which might induce local enterprises to hire. We emphasize the IV approach as the one that isolates the credit card market response to labor demand shocks.

V.B IV Results

Figure 4a shows the impact of a positive employment shock on accounts and credit limits. The corresponding coefficients are presented in Table II. We estimate that the number of accounts falls in response to a positive employment shock on impact, but longer differences attenuate this effect, with the relationship after four years no longer statistically distinguishable from zero. This pattern suggests important demand factors for short-term changes in the number of accounts. Figure 4b shows that credit limits also decline in response to a positive employment shock. These findings are consistent with the idea that when employment expands, consumers either apply for fewer accounts or close inactive accounts because of a lack of demand for unsecured credit.

We next examine how utilization measures of credit respond to employment. In Figure 5a, we show that log balances fall in response to a positive employment shock. A one percent increase in employment is associated with a 0.4 to 0.7 percent decline in balances, depending on the time horizon. Intuitively, as households receive positive employment shocks, they seek to pay down existing debts and are less likely to expand the amount of revolving credit.

Balances are the stock of outstanding debt, a function of purchases (inflows) and payments (outflows).[5] We find that both purchases and payments increase in response to a positive employ-

[5]Mechanically, our measure of the change in credit card balances is a function of balances across two periods, which are each a function of the purchases and payments made on the card. More directly, balances in period t are equal to balances in the previous period, minus any repayment, multiplied by $1+APR$, plus any additional purchases

ment shock, suggesting both increased consumption and repayment. These results are shown in Figures 6a and 6b. On net, given that balances decline, we estimate that the increases in payments are larger than the increases in purchases, but these differences are generally not statistically distinguishable from one another until a time horizon of at least 4 years.

V.C Heterogeneity

Our results thus far have focused on the average response to an employment shock. However, it is not necessarily the case that consumers would respond symmetrically to a positive employment shock vs. a negative shock. In particular, we might expect that negative shocks induce much larger and more rapid responses for credit demand.

It may also be the case that the effects are different for small vs. large employment changes. For instance, a large employment shock, such as a plant closing, may induce a widespread impact on the local economy and a larger shift in credit usage.

Finally, our data covers the period during the Great Recession and the slow recovery thereafter. The Great Recession was characterized by tremendous heterogeneity in the size of the downturn, based in part on exposure to the housing cycle (Mian and Sufi 2009). By comparing markets with larger vs. smaller housing cycles, we can examine whether the responsiveness of unsecured credit to employment shocks depended in part on the "debt overhang" of underwater mortgages.

V.D Robustness

The results thus far exploit county-level variation in employment shocks to examine credit market outcomes. By running long difference regressions, we use only two observations per county for each point-in-time difference (2008q1 relative to each subsequent quarter).

Alternatively, we can run quarterly panel specifications with a similar flavor. In this case, our instrument has more power because industrial composition is refreshed every period. However, this approach suffers from serial correlation. In Table IV, we show the quarterly panel estimation results. In short, many of the results are similar, but we find different patterns for the number of accounts across the two specifications. We are generally more skeptical of the quarterly panel

or fees.

results because of issues with serial correlation, which the long differences approach (which exploits one change per county) avoids.

V.E Interpretation of Magnitudes

Our empirics focus on estimating the elasticity of credit card variables with respect to employment, at the county level. Denote that elasticity for balances by

$$\epsilon = \frac{\frac{bal_{t+1}-bal_t}{bal_t}}{\frac{emp_{t+1}-emp_t}{emp_t}}$$

Suppose that the number of account-holders at time t equals the number of people with jobs. Then the balance per capita is $pcbal_t = bal_t/emp_t$. Now consider what happens when a single job loss occurs, i.e., $emp_{t+1} - emp_t = -1$. Then we can re-write the expression above as

$$\epsilon = -\frac{bal_{t+1} - bal_t}{pcbal_t}$$

The change in county level balances has two components: it reflects the change by the job-loser and the change by everyone else. If we assume that only the job loser adjusts, then the numerator equals the amount of his or her adjustment. If we assume a denominator (i.e., total credit card debt outstanding the period before a job loss) of approximately \$6000 (Ganong and Noel 2016), our estimated elasticity of $\epsilon = -0.5$ implies an increase in balances of \$3000 per marginal job-loser.

We can compare this figure to the reduction in income experienced upon job loss. Unemployment spells vary greatly in length, and grew on average during the Great Recession. Unemployment insurance replacement rates and durations also varied across states and time. A worker earning \$36000 per year who experienced a 90 day unemployment spell and received UI benefits with a replacement rate of 2/3 would have had an uninsured reduction in net income also equal to \$3000.

This suggests that the county-level aggregates could (under the above assumptions) result from approximately full self-insurance of shift-share-based income shocks with available credit card liquidity. However other factors unrelated to local labor demand, potentially including a national labor demand cycle, induced even larger shocks that credit card liquidity did not expand to smooth.

VI Conclusion

Our estimated responses to employment shocks, based on well-identified cross-sectional variation, provide a baseline for understanding how we would expect credit markets to respond to economic downturns. In particular, we estimate that negative employment shocks lead to increased credit card balances.

Notably, the aggregate patterns during our period of analysis go in the opposite direction. Large declines in employment nationwide during the Great Recession are associated with large decreases in balances. Figure 7 shows this relationship for Las Vegas, where the pattern is especially stark but not all that different from the relationship in the aggregate. While employment fell sharply by over 100,000 jobs, credit card balances fell by over 20 percent in Clark County.

How do we explain the discrepancy between our well-founded cross-sectional estimates and aggregate patterns that appear to go strongly in the opposite direction? The easiest explanation is that there were large supply-side credit shocks that occurred at the same time as the demand-side employment shocks we observe. This period was characterized by the collapse of housing values and the broader financial crisis.

The impact of tightening credit when demand was especially high had huge impacts on consumption during the Great Recession. While other work has focused on the "debt overhang" of secured debt such as mortgages, our results suggest that latent demand for unsecured credit likely far outstripped the amount of credit provided during the boom (Mian et al. 2013). Thus, our results complement those of Agarwal et al. (2016), who use a different identification approach, namely discontinuities in credit card offer algorithms to show the disconnect between credit supply and credit demand during the post-crisis recovery period.

Our estimates provide new evidence for the relationship between employment shocks and credit use. Sullivan (2008) and Keys (2017) explore the private-side responsiveness of credit markets to income shocks. More generally, our findings underscore concerns regarding unsecured credit access, which may not be available when needed most. The self-insurance aspects of credit cards are limited when issuers respond aggressively to deteriorating economic conditions, reducing credit lines and canceling accounts. Our results suggest that if credit supply amplifies business-cycle consumption volatility rather than mitigating it through consumption-smoothing, then other public and/or private sources must be considered to support the household sector beyond what would be expected in a static framework.

References

Agarwal, Sumit, Souphala Chomsisengphet, Neale Mahoney, and Johannes Stroebel, "Do Banks Pass Through Credit Expansions? The Marginal Profitability of Consumer Lending during the Great Recession," *Georgetown Working Paper*, 2016.

Aizer, Anna, "The Gender Wage Gap and Domestic Violence," *American Economic Review*, 2010, *100* (4), 1847.

Autor, David and Mark Duggan, "The Rise in the Disability Rolls and the Decline in Unemployment," *Quarterly Journal of Economics*, 2003, *118* (1), 157–206.

Bartik, Timothy J., *Who Benefits from State and Local Economic Development Policies?*, W.E. Upjohn Institute for Employment Research, 1991.

Blanchard, Olivier and Lawrence Katz, "Regional Evolutions," *Brookings Papers on Economic Activity*, 1992, *1992* (1), 1–75.

Bound, John and Harry Holzer, "Demand Shifts, Population Adjustments, and Labor Market Outcomes during the 1980s," *Journal of Labor Economics*, 2000, *18* (1), 20–54.

Chodorow-Reich, Gabriel, Laura Feiveson, Zachary Liscow, and William Gui Woolston, "Does State Fiscal Relief during Recessions Increase Employment? Evidence from the American Recovery and Reinvestment Act," *American Economic Journal: Economic Policy*, April 2012, *4* (3), 118–45.

Diamond, Rebecca, "The Determinants and Welfare Implications of US Workers' Diverging Location Choices by Skill: 1980–2000," *American Economic Review*, 2016, *106* (3), 479–524.

Ganong, Peter and Pascal Noel, "How Does Unemployment Affect Consumer Spending?," *Working Paper*, 2016.

Gorbachev, Olga and María José Luengo-Prado, "The Credit Card Debt Puzzle: The Role of Preferences, Credit Risk, and Financial Literacy," *Review of Economics and Statistics*, forthcoming.

Gould, Eric, Bruce Weinberg, and David Mustard, "Crime Rates and Local Labor Market Opportunities in the United States: 1979–1997," *Review of Economics and Statistics*, 2002, *84* (1), 45–61.

Greenstone, Michael, Alexandre Mas, and Hoai-Luu Nguyen, "Do Credit Market Shocks Affect the Real Economy? Quasi-Experimental Evidence from the Great Recession and 'Normal' Economic Times," *American Economic Journal: Economic Policy*, 2015.

Hoynes, Hilary, Douglas Miller, and Jessamyn Schaller, "Who Suffers During Recessions?," *Journal of Economic Perspectives*, 2012, *26* (3), 27–47.

Keys, Benjamin, "The Credit Market Consequences of Job Displacement," *Review of Economics and Statistics*, 2017.

—— **and Jialan Wang**, "Minimum Payments and Debt Paydown in Consumer Credit Cards," Technical Report, National Bureau of Economic Research 2016.

Maestas, Nicole, Kathleen Mullen, and David Powell, "The Effect of Population Aging on Economic Growth, the Labor Force and Productivity," *National Bureau of Economic Research Working Paper*, 2016, (22452).

Mian, Atif and Amir Sufi, "The Consequences of Mortgage Credit Expansion: Evidence from the US Mortgage Default Crisis," *Quarterly Journal of Economics*, 2009, *124* (4), 1449–1496.

____ **, Kamalesh Rao, and Amir Sufi**, "Household Balance Sheets, Consumption, and the Economic Slump," *Quarterly Journal of Economics*, 2013, p. qjt020.

Murphy, Kevin and Robert Topel, "The Evolution of Unemployment in The United States: 1968-1985," *National Bureau of Economic Research Macroeconomics Annual*, 1987, *2*, 11–58.

Notowidigdo, Matthew, "The Incidence of Local Labor Demand Shocks," *National Bureau of Economic Research Working Paper*, 2013, (17167).

Sullivan, James, "Borrowing During Unemployment: Unsecured Debt as a Safety Net," *Journal of Human Resources*, 2008, *43* (2), 383–412.

Table I: CCDB Summary Statistics

	2008q4	2013q4
Retail APR	14.03	13.62
Balance	$2,939	$2,683
Credit limit	$11,009	$9,029
Purchase volume	$1,355	$1,944
Payment	$1,598	$2,001
# of accounts (mil)	90	77

Note: This table presents summary statistics from the CFPB's Credit Card Database (CCDB). The first five rows show the mean values at the beginning and end years of our available sample, 2008q4 – 2013q4, using the fourth quarter in both cases for seasonal comparability. The final row shows the number of accounts used to construct these averages. See Section III for more detail on the sample construction.

Table II: Impact of Instrumented Employment Shocks on the Credit Card Market

	1 Year	2 Years	3 Years	4 Years	5 Years
Log Accounts	-0.24	-0.38	-0.49	-0.014	0.33
	(0.038)	(0.052)	(0.091)	(0.14)	(0.24)
Log Balances	-0.73	-0.49	-0.64	-0.41	-0.47
	(0.061)	(0.062)	(0.11)	(0.17)	(0.3)
Log Limits	-0.37	-0.66	-1.04	-0.99	-1.37
	(0.044)	(0.064)	(0.12)	(0.2)	(0.4)
Purchase Volume	0.074	0.16	0.14	0.47	0.54
	(0.084)	(0.071)	(0.11)	(0.15)	(0.26)
Payment Amount	0.13	0.23	0.23	0.83	1.27
	(0.076)	(0.066)	(0.1)	(0.15)	(0.26)
APR	-0.083	-0.044	0.019	0.44	0.82
	(0.027)	(0.034)	(0.044)	(0.078)	(0.17)

Note: This table presents IV estimates based on long-difference specifications relating local credit card market features with emlpoyment shocks. Standard errors, clustered by county, are in parentheses. The long difference in each specification is benchmarked to 2008q1, so the one year column measures outcomes between 2008q1 and 2009q1, etc.

Table III: First Stage Relationship

Outcome Variable	First Stage County	First Stage State
Instrumented Employment	0.51 (0.016)	0.20 (0.079)

Note: This table presents the relationship between employment and instrumented employment based on industry-level shifts and pre-shock composition. Regressions include seasonal geography dummies and are weighted based on population. Standard errors clustered on geography are shown in parentheses. The first-stage F-statistics are 1168 for the county specification and 189 for the state specification.

Table IV: Responses to Employment Shocks in the Credit Card Market

	Employment	
Outcome Variable	OLS	IV
Log Accounts	0.027	0.15
	(0.0027)	(0.023)
Log Balances / Account	-0.047	-0.17
	(0.0034)	(0.029)
Log Limits / Account	-0.01	0.053
	(0.0014)	(0.012)
Purchase Volume / Account	0.11	0.21
	(0.0064)	(0.054)
Payment Amount / Account	0.11	0.26
	(0.0067)	(0.057)
APR	0.0068	-0.11
	(0.0025)	(0.022)

Note: This table uses the quarterly panel version of the data to estimate the responsiveness of the credit card market to employment shocks. Standard errors clustered at the county level are shown in parentheses. All specifications include county and time fixed effects, and estimates are weighted based on the number of credit card accounts.

Figure 1: Employment and Credit Card Balances over the Last Cycle

(a) Total Employment (in millions)

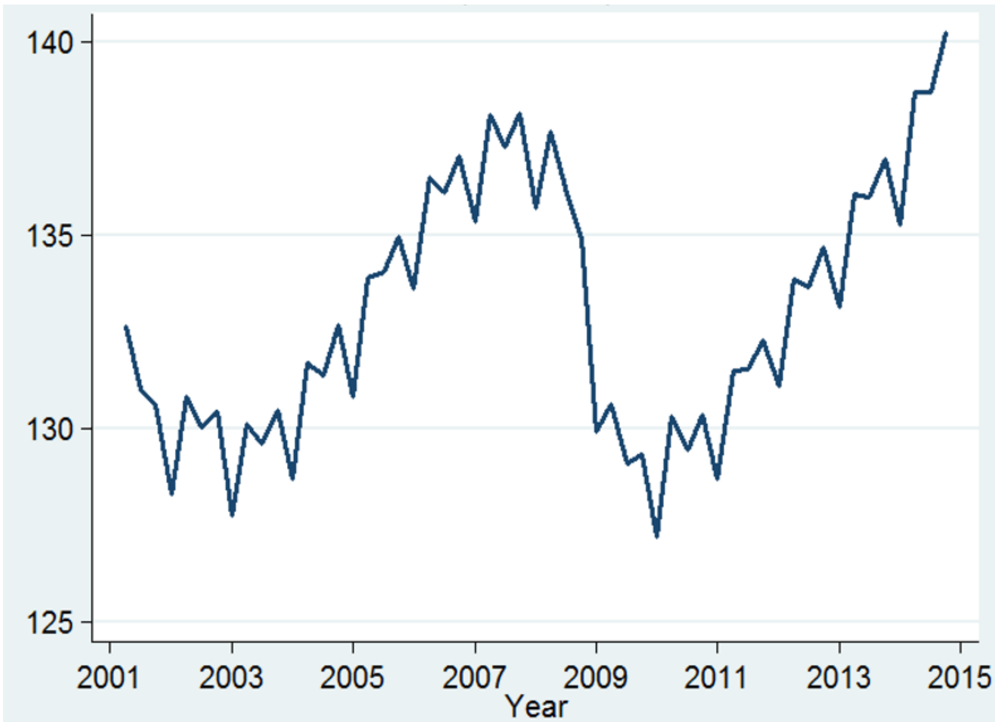

(b) Aggregate Credit Card Balance Outstanding (in trillions)

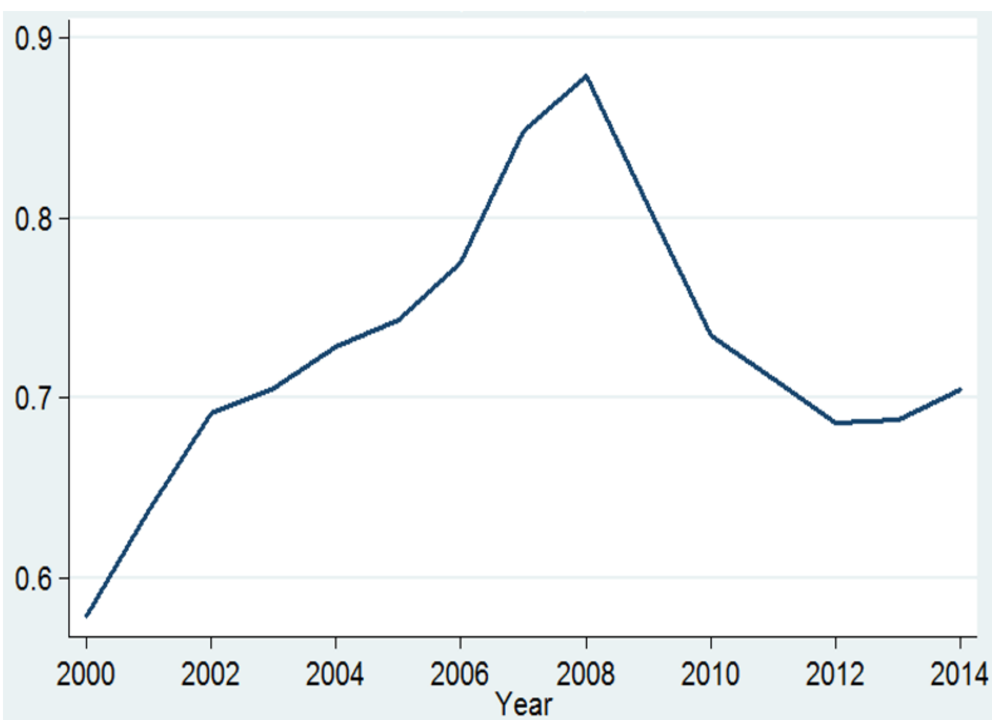

Sources: Bureau of Labor Statistics (panel a), Federal Reserve Bank of New York calculations based on Equifax/Consumer Credit Panel data.

Figure 2: First Stage Regression, County-by-Quarter Level

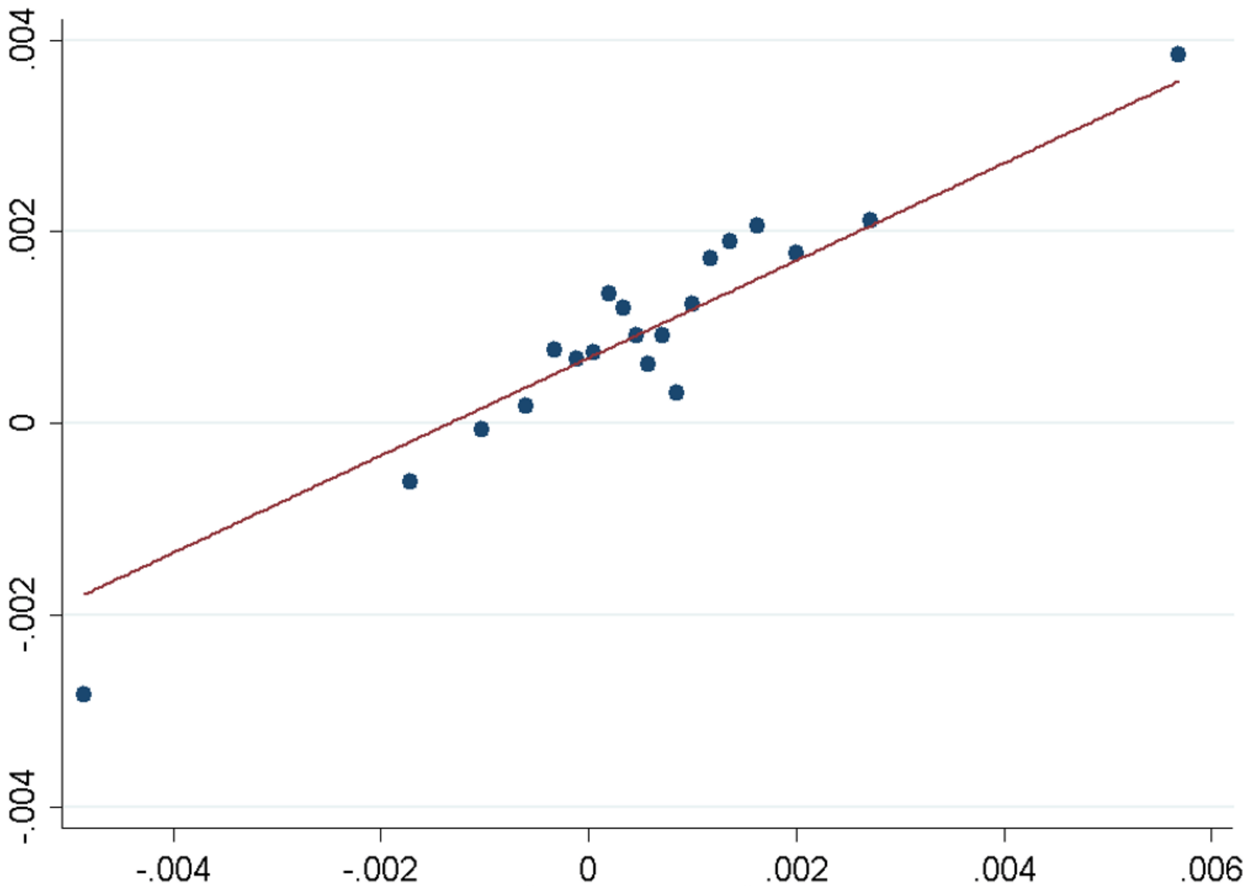

Note: This figure shows the bin-scatter relationship between employment and instrumented employment. The figure presents the relationship using bins of the x-axis. The original data is based on all counties and all quarters from 2008q1 to 2013q4.

Figure 3: First Stage Long-Difference Specifications

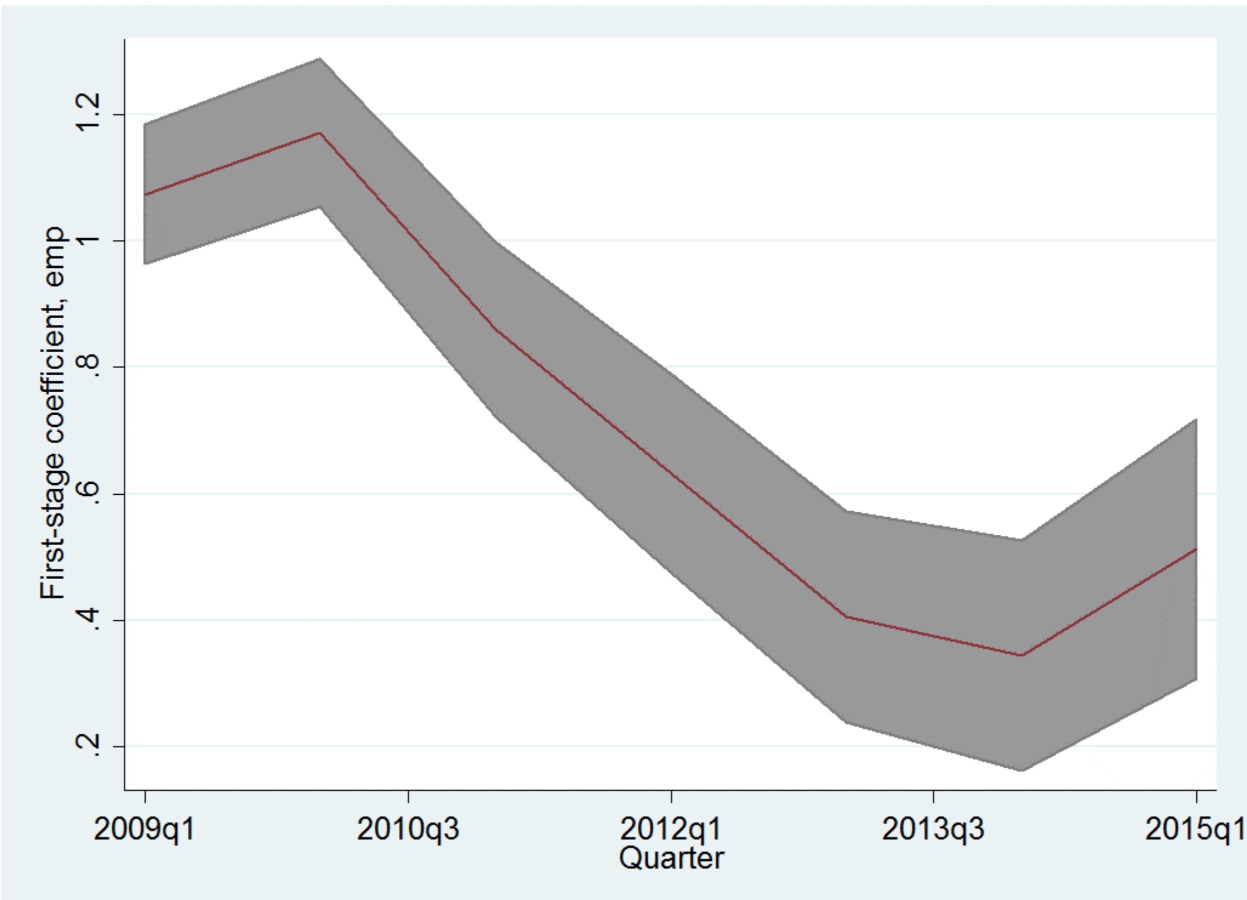

Note: This figure presents the relationship between employment and instrumented employment in the long-difference specification. Each data point is estimated using a separate regression, quarter-by-quarter, and the gray bars represent 95% confidence intervals.

Figure 4: Long-Difference Relationship between Employment and Accounts, Credit Limits

(a) Employment vs. Accounts

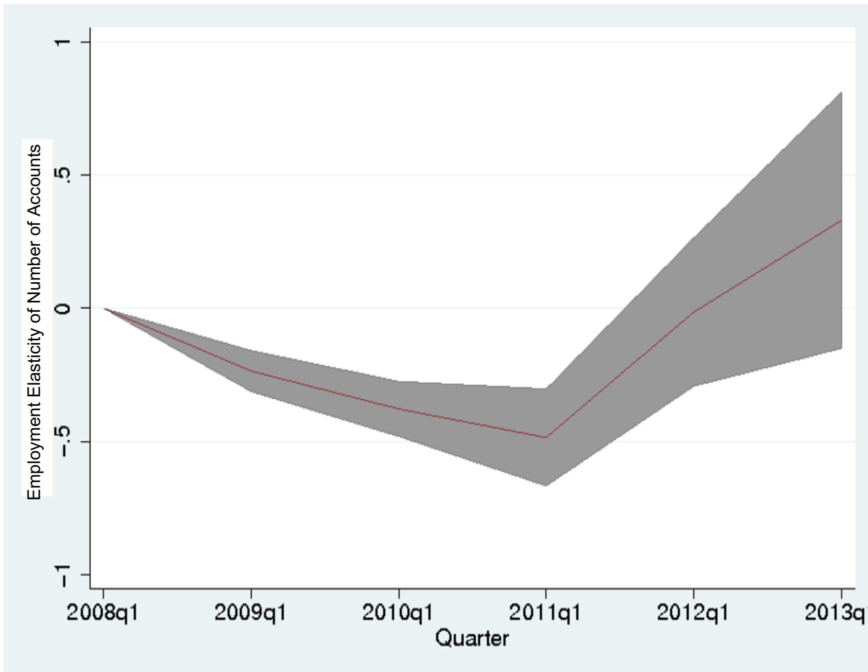

(b) Employment vs. Credit Limits

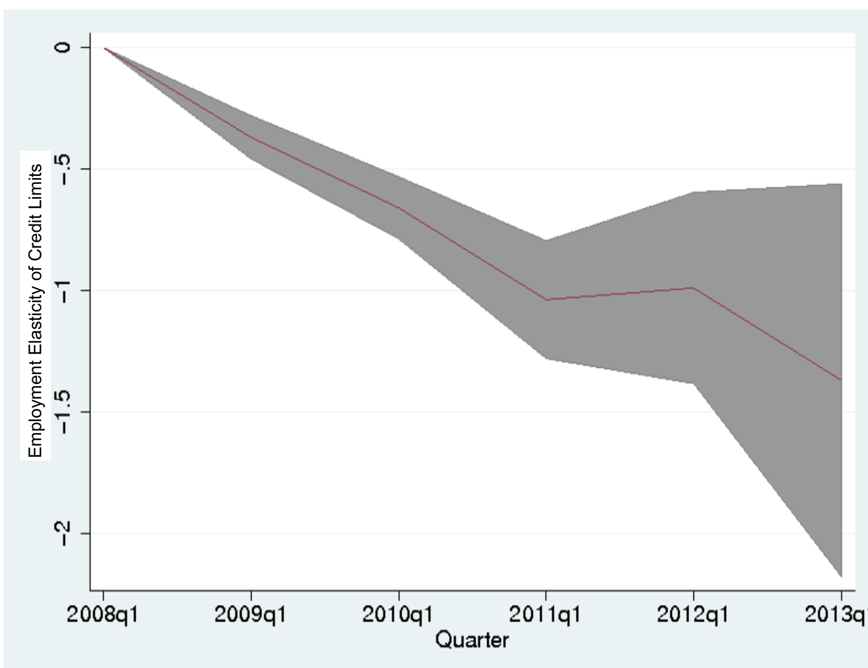

Note: The top panel of this figure presents the relationship between the number of accounts and instrumented employment in the long-difference specification. The bottom panel presents the relationship between average credit limits and instrumented employment. Each data point is estimated using a separate regression, quarter-by-quarter, and the gray bars represent 95% confidence intervals.

Figure 5: IV Relationship over Time

(a) Employment vs. Credit Card Balances

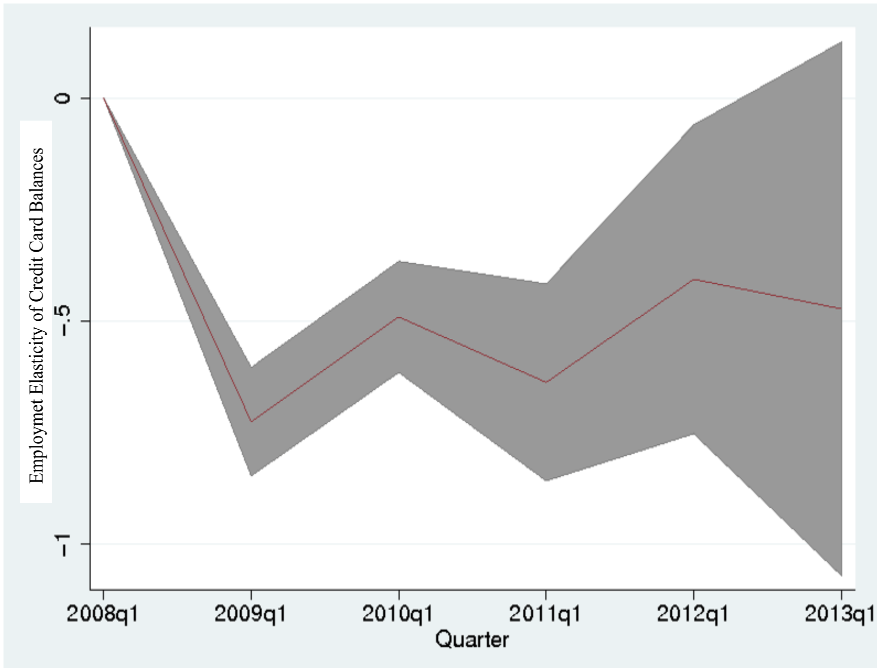

(b) Employment vs. APR

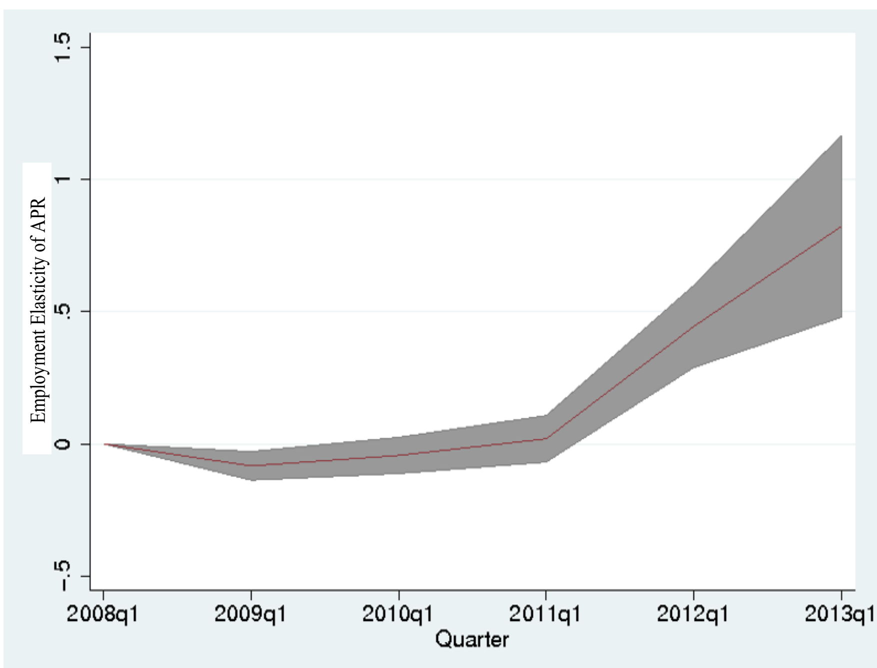

Note: The top panel of this figure presents the relationship between credit card balances and instrumented employment in the long-difference specification. The bottom panel presents the relationship between average APR and instrumented employment. Each data point is estimated using a separate regression, quarter-by-quarter, and the gray bars represent 95% confidence intervals.

Figure 6: IV Relationship over Time

(a) Employment vs. Credit Card Purchases

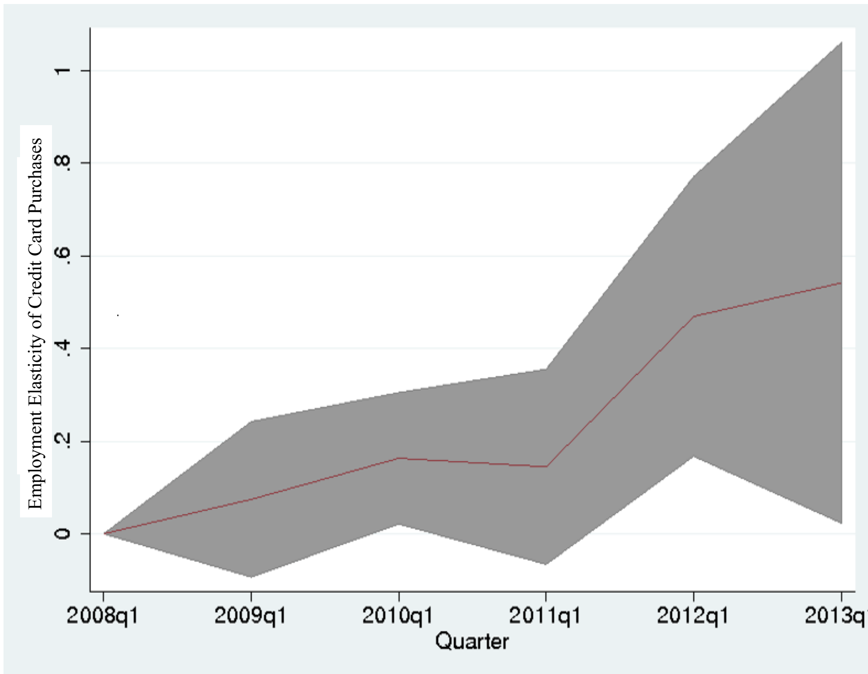

(b) Employment vs. Payments

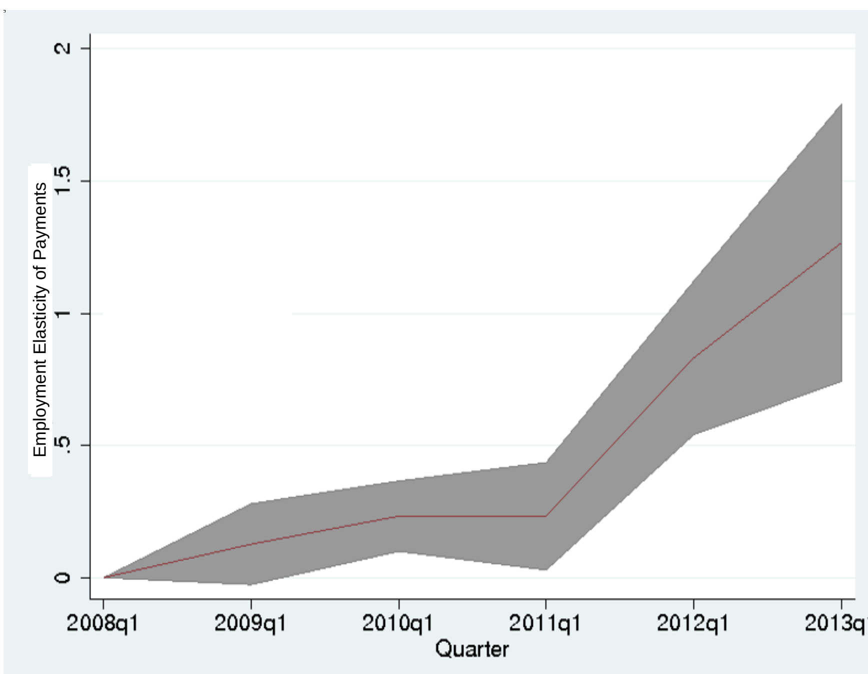

Note: The top panel of this figure presents the relationship between credit card purchases and instrumented employment in the long-difference specification. The bottom panel presents the relationship between credit card payments and instrumented employment. Each data point is estimated using a separate regression, quarter-by-quarter, and the gray bars represent 95% confidence intervals.

Figure 7: Time Series for Clark County, NV

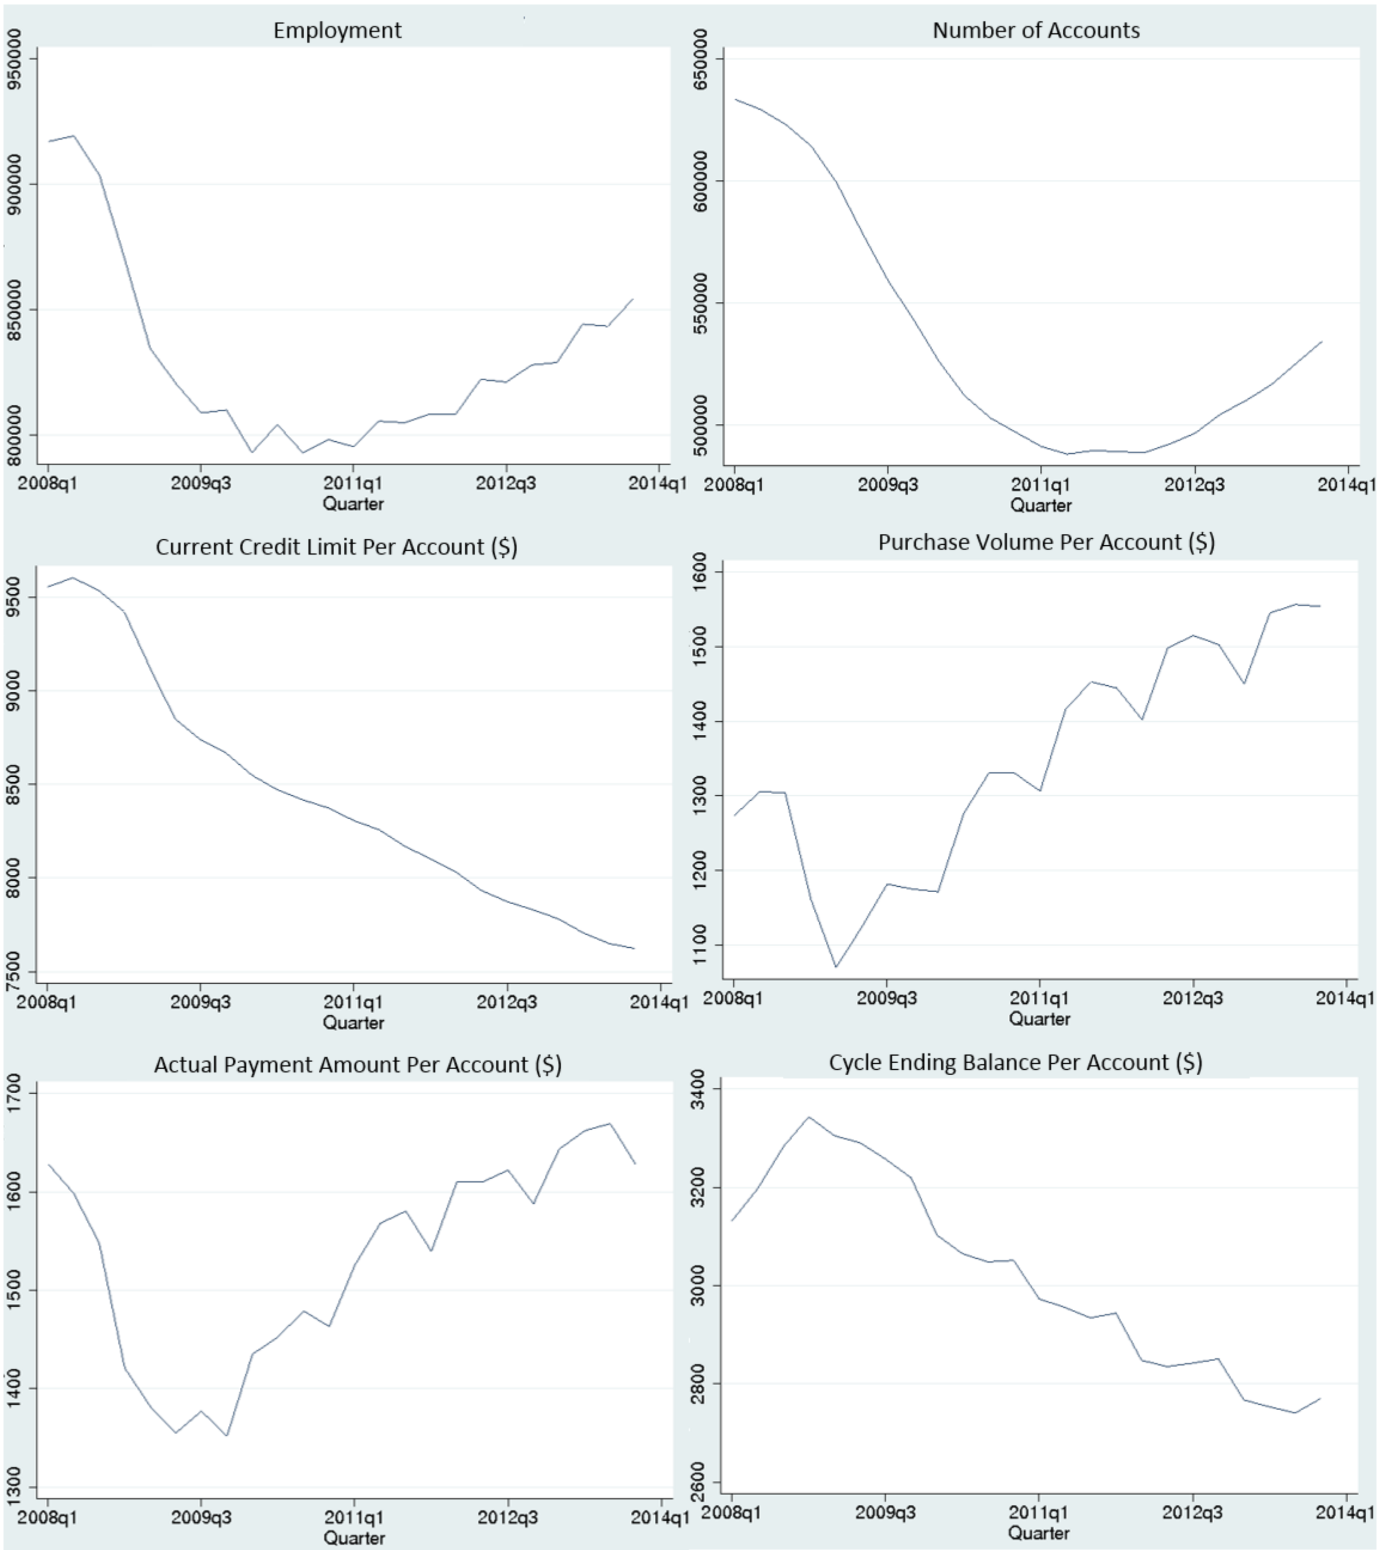

Note: This figure presents the time series of employment and various credit card attributes for Clark County, NV (home of Las Vegas) from 2008q1 to 2013q4.

Figure 8: Responses to Employment Shocks in the Credit Card Market

(a) Number of Accounts

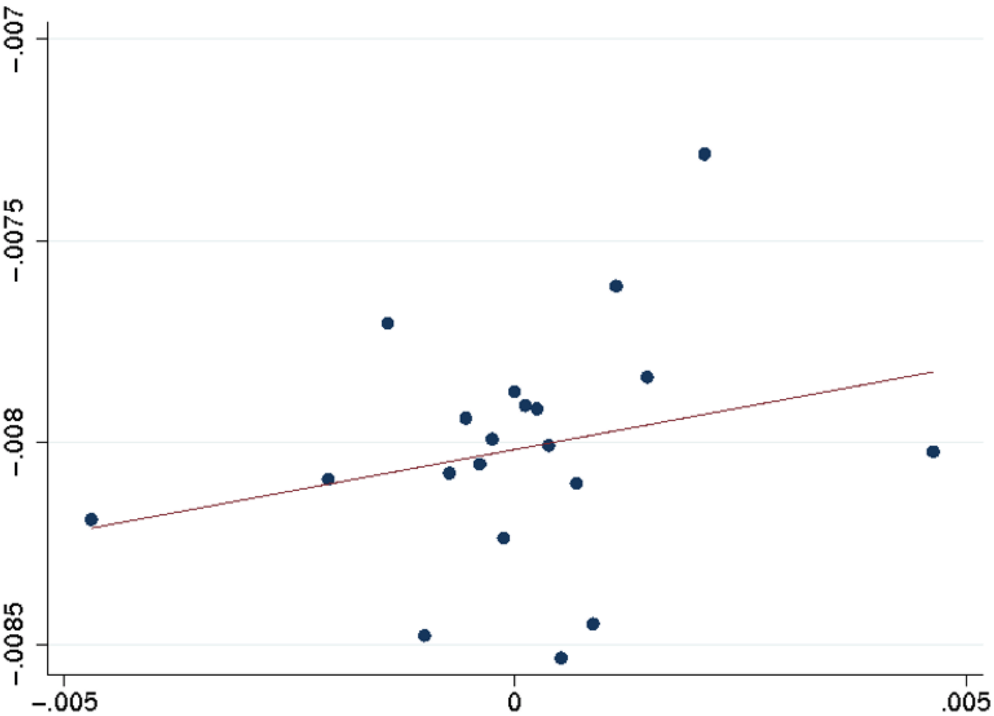

(b) APR

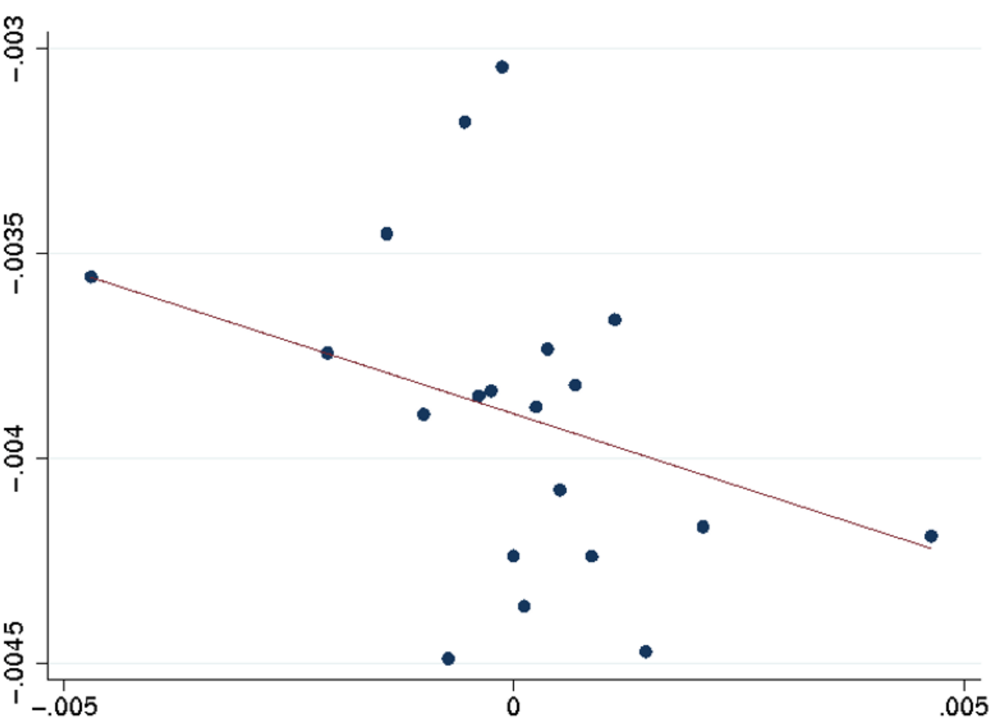

Note: This figure shows the relationship between card attributes, namely the number of accounts (panel a) and average APR (panel b), and instrumented employment shocks. The original data is based on county-quarter level data from 2008q1 to 2013q4.

Figure 9: Impact of Employment Shock on Credit Card Balances

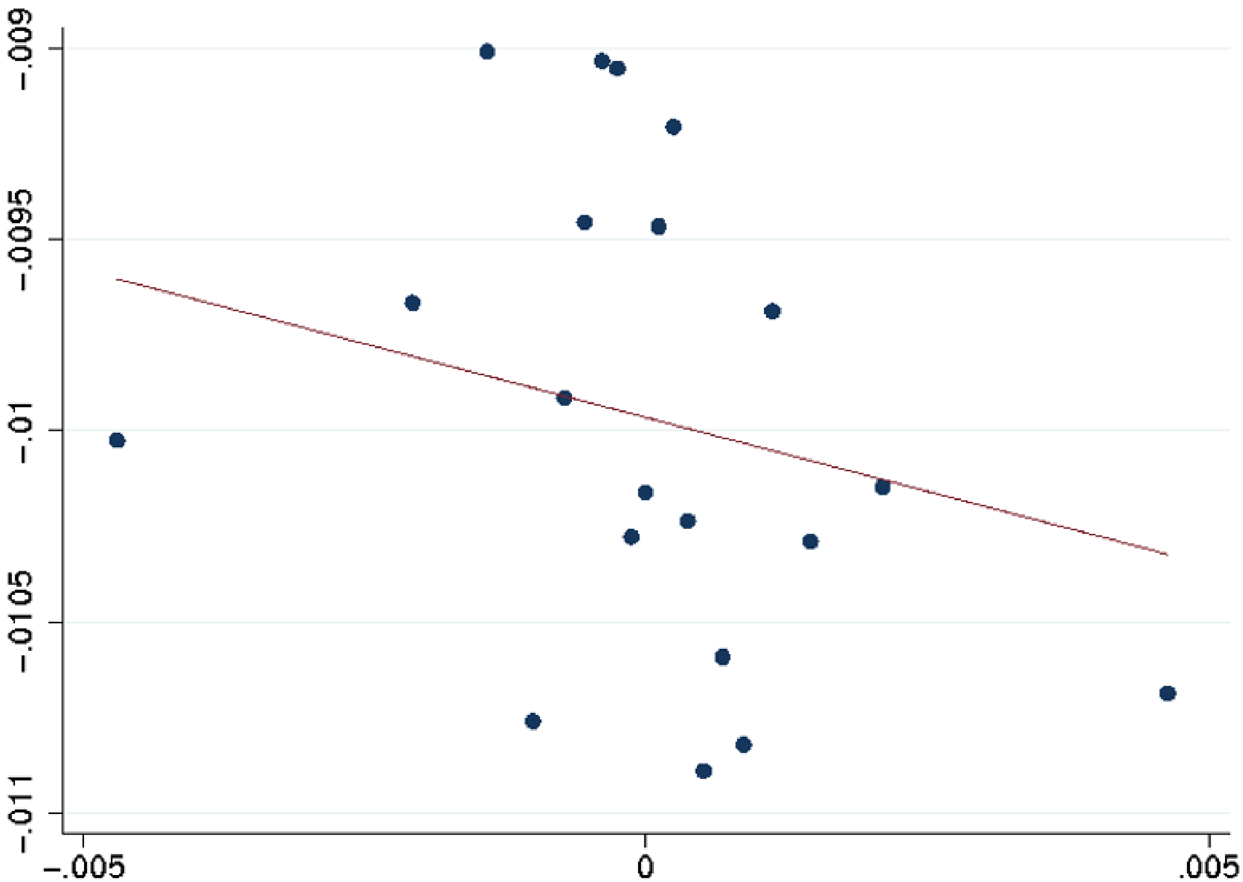

Note: This figure shows the relationship between average credit card balances and instrumented employment shocks. The original data is based on county-quarter level data from 2008q1 to 2013q4.

Figure 10: Dispersion Across Counties

(a) Credit Limit per Account

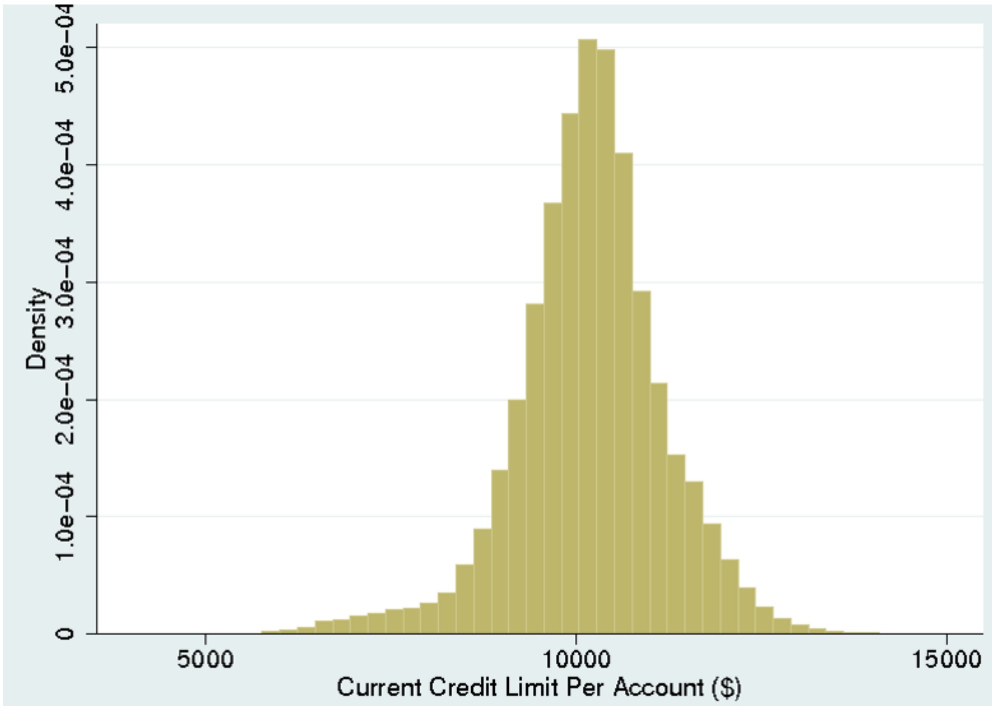

(b) Average APR

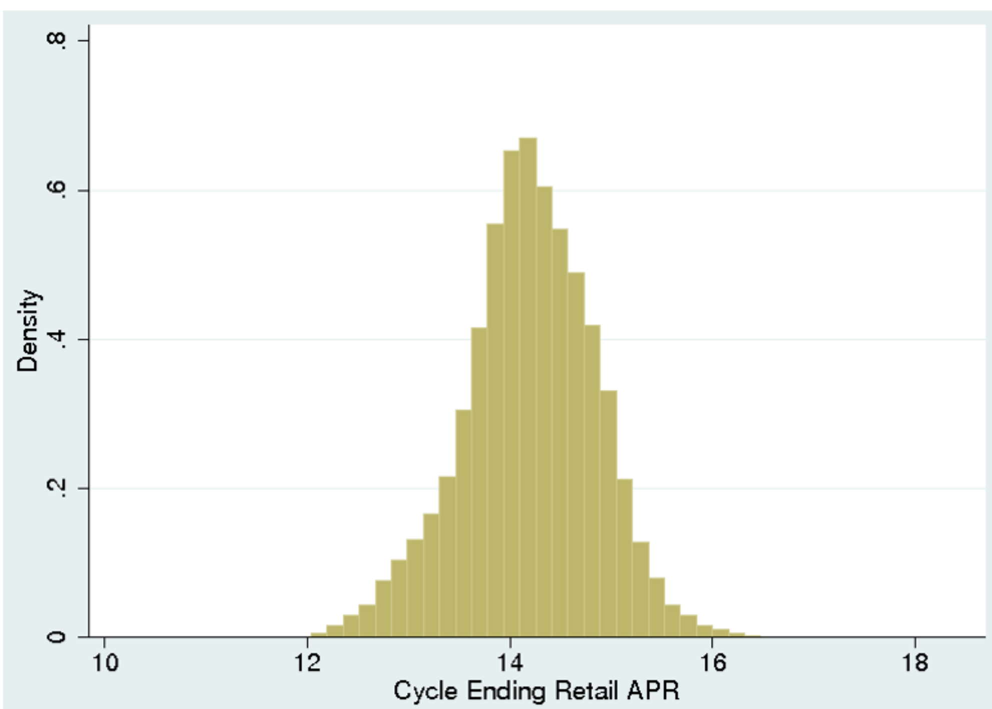

Note: This figure presents the distribution of average credit limits per account (panel a) and average APRs (panel b), by county, in 2008q1.

Figure 11: Employment and Credit Card Accounts in Clark County, NV

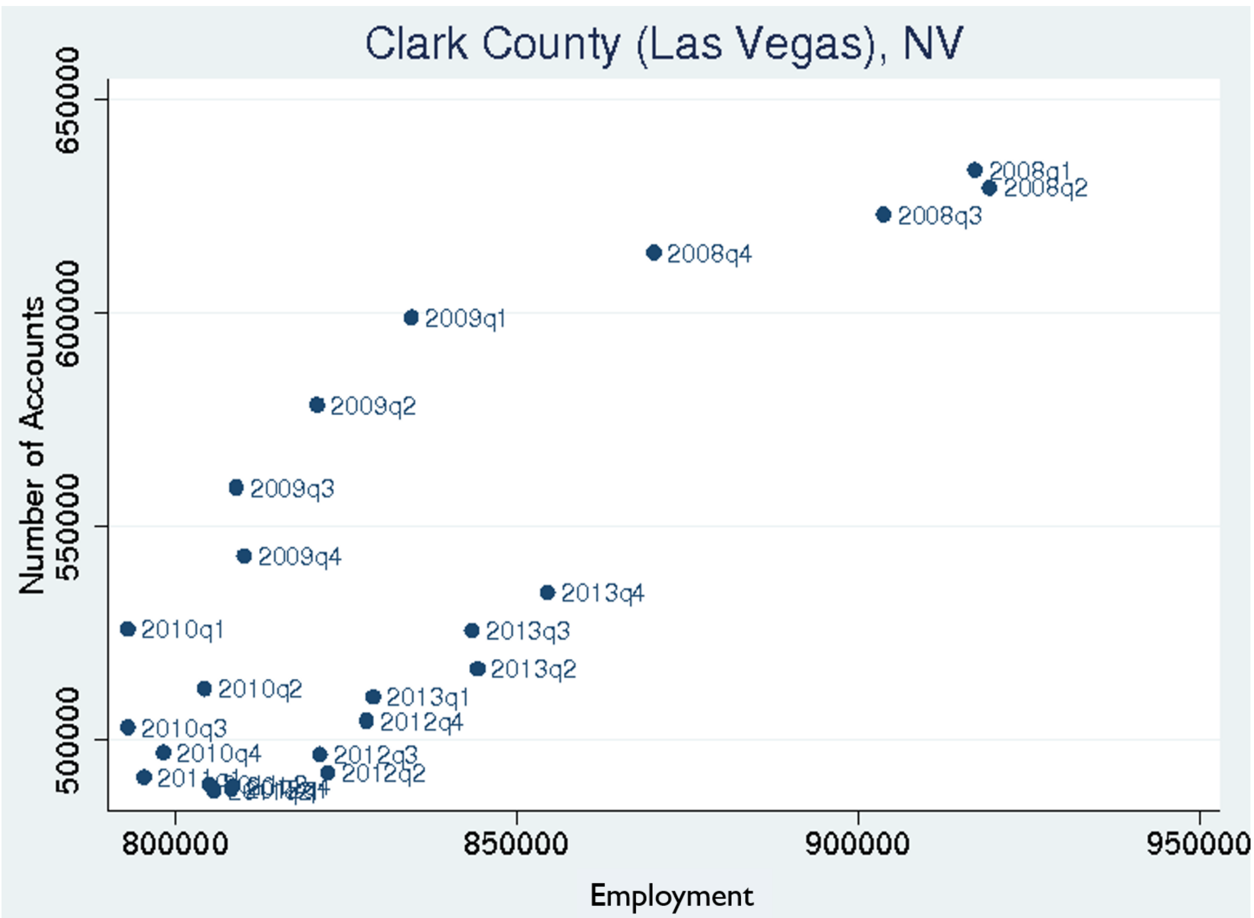

Note: This figure shows the relationship between employment and the number of credit card accounts quarter-by-quarter from 2008q1 to 2013q4.