

The Decline in Currency Use at a National Retail Chain

Zhu Wang and Alexander L. Wolman

The composition of US retail payments is changing rapidly. According to the Federal Reserve's triennial Payments Study (2013, 2016), from 2012 to 2015 the value of debit and credit card payments increased at annual rates of 7.1 percent and 7.4 percent, respectively. Over this same period, nominal GDP rose at less than a 4 percent annual rate, which suggests that the increase in card payments came at the expense of some other form(s) of payments, the obvious candidates being checks and cash. The value of check payments did fall over this period, but it is possible that the fall in check payments was offset by an increase in ACH rather than card payments; ACH tends to be used in business and financial transactions while cards are used in consumer payments. The Payments Study covers only noncash payments, but Wang and Wolman (2016a) provide direct evidence about cash use at a large discount retailer, finding that the cash share of the number of payments fell by 2.46 percentage points per year from 2010 to 2013. In their study, an increase in card use was almost the mirror image of a decrease in cash use.

At least four sets of factors could be contributing to the apparent shift from cash to card in retail payments. First, Wang and Wolman (2016a) documented a negative relationship between transaction size and the share of cash transactions; thus, some of the decline in observed cash shares could be due to an increase in average transaction size. Second, Wang and Wolman also documented systematic relationships between the cash share of payments in a location and the

■ The views in this paper are those of the authors and do not necessarily represent the views of the Federal Reserve Bank of Richmond, the Federal Reserve Board of Governors, or the Federal Reserve System. The authors thank Erica Paulos for excellent research assistance, and Mohamed Abbas Roshanali, Arantxa Jarque, Bruno Sultanum, and John Weinberg for helpful comments on an earlier draft.

demographic and economic characteristics of the location; over time, changes in those characteristics may explain changes in the cash share. Third, changes in technology may be reducing the cost and increasing the availability and security of debit and credit cards. And fourth, consumers' perceptions of cards may be improving slowly, generating a gradual expansion in card use. This paper brings new evidence to bear on the contributions of the first two factors to the decline in cash payments. Using an updated version of the data from Wang and Wolman (2016a,b), we study the association between changes in payment shares and changes in the size of transactions as well as changes in location-specific economic and demographic variables over the period from February 2011 to February 2015. While we cannot distinguish the third and fourth factors listed above, the portion of the decline in cash shares that is unexplained by our analysis represents the sum of these two sets of factors.

There are important public policy questions for which it matters what explains the decrease in cash use. Cash remains an important means of payment in the United States, and in the wake of the long recent experience with interest rates at their effective lower bound, some economists have advocated policies that would reduce or even eliminate the availability of paper currency (Rogoff 2016). Without paper currency, the argument goes, monetary policy would no longer be constrained by a lower bound on nominal interest rates.¹ Against this, the benefits of cash must be considered, and the accounting we provide for the decline in cash use can contribute to the debate over the benefits of cash. To the extent that the decline in cash use is accounted for by changing demographics or changing transaction size, there may be greater scope for concern about the effects of a (hypothetical) elimination of currency on particular segments of society.

In Wang and Wolman (2016a), and in this paper, we analyze transactions data from a discount retailer with thousands of stores across the US. In the earlier paper, we combined the transactions data with fixed demographic data and other data across locations.² With almost two million transactions every day, we were able to precisely characterize the daily and weekly patterns of payment use. And, with thousands of zip-code locations, we were also able to precisely estimate the relationships between cash shares and location-specific variables. However, the fact that our data covered only three years meant that we could not

¹ Rogoff (2016) also sees benefits from eliminating cash related to the fact that cash is heavily used in the underground economy.

² In Wang and Wolman (2016b), we conducted a similar analysis that concentrated on retail outlets in the Fifth Federal Reserve District.

incorporate time variation in the location-specific data: the Census Bureau’s American Community Survey (ACS) data were not available at the zip-code level for more than one year in our dataset. In the current paper, we do not attempt to capture the daily variation in payment shares but instead focus on the “medium-term” shift in the cash share of transactions from February 2011 to February 2015, using only data from those two months. While we sacrifice on one dimension, we are able to incorporate time variation in the location-specific data using the five-year ACS estimates at the zip-code level for 2011 and 2015.

On average, across the stores in our study, the share of cash transactions fell by 8.6 percentage points from February 2011 to February 2015. Our statistical model attributes approximately 1.3 percentage points of that decline to increasing transaction sizes. Changes in demographic and other location-specific variables contribute between 0.5 and 1.3 percentage points, so our analysis attributes approximately three-quarters of the decline in cash use to a pure time effect, which stands in for the third and fourth factors listed above, and any other factors omitted from our analysis.

1. TRANSACTIONS DATA: THE DECLINE IN CURRENCY USE

Our payments data come from a US retail chain selling a wide variety of goods, with a majority of its revenue accounted for by household consumables such as food and health-and-beauty aides. The chain has thousands of stores and is located in most states. Although there is not a specific geographic focus, the stores tend to be located in relatively low-income zip codes.³ While the raw data are at the level of individual transactions (time and location, size, means of payment), our analysis uses aggregated data: for each zip code, we compare the shares of transactions in each of the four main payment types (cash, debit card, credit card, and check) in February 2011 to the corresponding shares in February 2015. One month is a long enough time period to get a relatively large number of transactions: most zip codes had more than 7,000 transactions in each of the two months. The total number of zip-code locations is more than 5,000. We chose February 2011 and February 2015 to balance two considerations. A longer time span provides a better sense of the trend decrease in cash use, but we needed

³ See Wang and Wolman (2016a) for some additional information. Our use of the data is governed by a confidentiality agreement that limits the degree of detail we may disclose.


to choose years for which zip-code-level data are available from the ACS.

Figure 1 is a scatterplot of the share of cash transactions in each zip code in 2015 and 2011, on the y- and x-axes respectively. The solid grey line is the locus of points for which the cash share is equal in the two years, and points below (above) the line indicate a decrease (increase) in the cash share. This figure provides a nice overview of the data and the properties we want to study. First, there is significant variation in the share of cash transactions in both years. Second, the share of cash transactions declined from 2011 to 2015 in almost every zip code, as indicated by the small number of observations that lie above the $y=x$ line. And third, while the decrease in the cash share does not seem closely related to the level of the cash share, the decrease is also not constant across zip codes. The first and third properties—cross-zip-code variation in both the level and change in the cash share—provide motivation for using demographic and other zip-code-level variables in our statistical analysis. The second property—a significant common component in the change in the cash share across zip codes—could partly reflect changes in demographics that are common across locations. However, the common component also reflects changes in payments technology and consumer perceptions that are not captured by our analysis.

Table 1 displays summary statistics for the data in Figure 1, as well as the corresponding data for shares of debit, credit, and check transactions. From February 2011 to February 2015, the average cash share of transactions across zip codes declined from 78.2 percent to 69.5 percent, or 2.18 percentage points per year. Our focus is primarily on the decline in cash and the combined increase in credit and debit use; the total card share of transactions increased by an average of 2.3 percentage points per year, with the difference, 0.12 percentage points per year, accounted for by a decrease in the share of transactions conducted with checks. Our data are not well-suited to distinguishing credit and debit transactions because the category we call “debit” includes only PIN debit transactions—signature debit and most prepaid cards are included in “credit.”⁴ PIN debit transactions increased by an average of 1.63 percentage points per year, approximately 70 percent of the overall increase in card use.

Table 1 also shows that from 2011 to 2015 both the standard deviation of cash transaction shares and the interquartile range (difference between the 75th and 25th percentiles) increased. This corresponds to

⁴ PIN debit is a debit card transaction that requires the consumer to enter a PIN number, whereas signature debit is a debit card transaction that requires the consumer to sign their name (like a credit card transaction).

Figure 1 Zip-Code-Level Cash Shares for 2015 and 2011

the third property noted in reference to Figure 1: the distribution of cash shares across zip codes did not shift down in a uniform manner. Figure 2 illustrates this explicitly, showing that the histogram of cash shares across zip codes was more spread out in 2015 than in 2011, in addition to shifting to the left.

Dispersion across locations in the *change* in cash shares is illustrated in the third row of Table 1 and in Figure 3. Cash shares declined by an average of 8.6 percentage points, but there is significant dispersion: in 25 percent of zip codes, the cash share decreased by at least 9.9 percentage points, and in 25 percent of zip codes the cash share decreased by less than 7.0 percentage points.

As mentioned in the introduction, one factor that could help account for the changes in cash shares depicted in Figures 1 through 3 is a change in the distribution of transaction sizes. Our econometric analysis of the change in cash shares below will explicitly take into account transaction size, but for now we simply report on the distributions of median transaction size and change in median transaction size by location. Table 2 provides various statistics for the distributions: for example, the mean value of median transaction size rose from \$7.26 to \$7.96, and the mean change in median transaction size is \$0.70.

Table 1 Payment Shares Across Zip Codes, February 2011 vs. February 2015

	Mean	Std. dev.	1%	25%	50%	75%	99%
Cash: 2011	0.782	0.056	0.636	0.747	0.787	0.822	0.891
2015	0.695	0.063	0.532	0.653	0.699	0.740	0.824
change	-0.086	0.025	-0.150	-0.099	-0.085	-0.070	-0.031
Debit: 2011	0.161	0.050	0.062	0.127	0.156	0.192	0.292
2015	0.226	0.058	0.095	0.187	0.222	0.261	0.380
change	0.064	0.028	-0.016	0.049	0.065	0.081	0.128
Credit: 2011	0.047	0.034	0.008	0.024	0.036	0.059	0.171
2015	0.074	0.049	0.015	0.039	0.060	0.096	0.246
change	0.027	0.029	-0.017	0.009	0.019	0.039	0.121
Check: 2011	0.010	0.011	0.000	0.002	0.006	0.014	0.051
2015	0.005	0.006	0.000	0.001	0.003	0.007	0.026
change	-0.006	0.006	-0.027	-0.008	-0.004	-0.001	0.001


Note: Rows titled “change” show distributions of changes in payment shares from 2011 to 2015. These may show different means than the change in the mean share for a particular payment type because the set of stores is not identical in the two years (e.g., for cash, change in mean is 0.087 and mean change is 0.086).

Table 2 Median Size of Transactions Across Zip Codes, February 2011 vs. February 2015

	Mean	Std. dev.	1%	25%	50%	75%	99%
2011	7.26	1.02	5.35	6.56	7.15	7.81	10.12
2015	7.96	1.10	5.87	7.20	7.88	8.66	10.90
Change	0.70	0.78	-1.28	0.27	0.67	1.12	2.52

Note: The third row is the distribution of change in median transaction size from 2011 to 2015.


Figures 4 and 5 display histograms of the two distributions of median transaction size (Figure 4) and the distribution of changes in median transaction size (Figure 5). The distribution of transaction sizes shifted to the right from 2011 to 2015 and became slightly more spread out. The dispersion in changes in median transaction size (Figure 5) is indeed consistent with the behavior of transaction size accounting for some of the shift in the cash share distribution from 2011 to 2015.

Figure 2 Histograms of Zip-Code-Level Cash Share

2. LOCATION-SPECIFIC DATA

Table 3 provides summary statistics for the location-specific data used in our analysis, comparing the 2011 and 2015 values. Wang and Wolman (2016a) provide a discussion of why one would expect these variables to be relevant for explaining payment choice, arguing that each consumer has a threshold transaction size below which they will use cash and above which they will use a noncash form of payment. The threshold may vary over the week, month, and year, and it will likely be related to the consumer's financial situation, their demographic characteristics, and their surrounding environment (including banking options, population density, and crime rates). The overall cash share in a particular location at a particular time will thus depend on the characteristics of the consumers in that location, the characteristics of the location, and the size distribution of transactions.

In Wang and Wolman (2016a), we used the same demographic variables to account for variation in cash shares across locations, but our data did not allow for the possibility of using changes in those variables to account for the change over time in cash shares; the location-specific variables were necessarily treated as fixed over the three-year sample of data due to limitations of the Census Bureau data. Here, the longer


Figure 3 Histogram of Change in Zip-Code-Level Cash Share

span of the transactions data means we can incorporate distinct demographic data for 2011 and 2015 for each zip code to decompose the changes in cash shares. Our earlier paper used forecasted nationwide changes in the location-specific variables to project *future* changes in cash shares and attributed up to 15 percent of the overall projected decline in cash shares to forecasted changes in location-specific variables. Below, we will compare that number to our decomposition of actual changes in cash shares.

The demographic variables (sex, age, race, and education) and the housing variables in Table 3 are all from the ACS. We use ACS five-year estimates at the zip-code level for 2011 and 2015. Note that for age we report only the 2011 data. We fix the age data at 2011 levels because we think that cohort is more important than age for payment behavior.⁵ The banking variables—market concentration, as measured by the

⁵ In principle, we would like to use data on the distribution of cohorts in each year. However, because the age data in our regression are in relatively large bins (e.g., fifteen years), it will not provide an accurate picture of how the cohort distribution changes across the four-year span of our data. In Section 4, we will use the estimated coefficients together with more detailed age data to construct a rough measure of the cohort contribution to the change in cash shares.


Figure 4 Histogram of Zip-Code-Level Median Transaction Size



Herfindahl-Hirschman index (HHI), and the number of bank branches per capita—are from the FDIC’s Summary of Deposits. Banking HHI is calculated by squaring each bank’s share of deposits in a zip code and then summing these squared shares. We allow the HHI effect to differ between rural and urban areas because of the possibility that high concentration in an urban area may reflect the presence of a small number of high-productivity banks. The robbery rate is from the FBI’s uniform crime report (note that the robbery rate is at the county level). In most cases, the changes from 2011 to 2015 appear to be small.⁶ However, the examples of median household income and education show that changes in location-specific variables have the potential to account for some of the decline in cash use. Across locations, Wang and Wolman (2016a) found that higher educational attainment and higher income were associated with lower cash use; Table 3 shows that both educa-

⁶ One exception is the HHI index. Note that in our earlier work the HHI was measured at the level of metropolitan statistical area (MSA) or rural county. Here it is measured at the zip-code level. In Wang and Wolman (2016b), we found that variation in HHI explained little of the variation in payment shares across zip codes.

Figure 5 Histogram of Change in Zip-Code-Level Median Transaction Size



tional attainment and income increased on average from 2011 to 2015, which would be consistent with a decrease in cash use assuming the relationship found by Wang and Wolman also holds across time. In the next section, we will report estimates of a statistical model similar to that in our 2016a paper using the variables in Table 3. Then in Section 4, we will quantify the contributions of changes in transaction size and in the demographic variables to the decline in cash use.

3. EMPIRICAL FRAMEWORK AND ESTIMATES

In this section, we describe the statistical model used to analyze payment shares and provide a summary of the estimates. The statistical model is tailored to the properties of the variable we are seeking to explain: in a particular time period in a particular location, the shares of cash and other payment types are each between zero and one, and they must sum to one. These properties mean that linear regression is not appropriate.

Table 3 Summary Statistics of Zip-Code Variables

Variable (unit)	Mean		Std. dev.	
	2011	2015	2011	2015
Banking HHI	0.43	0.46	0.26	0.26
Banking HHI \times Metro	0.28	0.29	0.29	0.30
Branches per capita ($1/10^3$)	0.38	0.36	0.36	0.32
Robbery rate ($1/10^5$)	13.17	12.34	28.477	26.02
Median household income (\$)	43,221	43,818	12,289	12,621
Population density (per mile ²)	1479	1484	2614	2493
Family households (%)	66.50	65.52	8.65	8.85
Housing: Renter-occupied (%)	28.18	30.14	11.21	11.79
Owner-occupied	57.33	55.28	12.86	12.77
Vacant	14.49	14.58	8.59	8.63
Female (%)	50.87	50.74	2.87	2.92
Age: < 15 (%)	20.03	-	4.08	-
15-34	26.65	-	5.88	-
35-54	27.36	-	3.28	-
55-69	16.16	-	3.77	-
≥ 70	9.81	-	3.81	-
Race: white (%)	74.88	75.62	22.80	22.18
black	16.61	15.85	21.65	20.94
Hispanic	13.55	15.26	19.39	20.83
Native	1.07	1.06	4.20	4.08
Asian	1.42	1.58	2.34	2.61
Pac-Islr	0.06	0.06	0.28	0.30
other	3.81	3.31	6.31	5.36
multiple	2.15	2.51	1.76	1.92
Educ below high school (%)	18.36	16.89	8.70	8.61
high school	34.22	33.62	7.33	7.41
some college	21.28	21.76	4.34	4.21
college	26.14	27.72	10.18	10.47

Note: The sum for race percentage is greater than 100 because Hispanic includes other categories.

Description of model

The purpose of the statistical model is to provide estimates of the relationship between the levels of payment shares and a set of explanatory variables comprising transaction size, the time- and location-specific variables, state-level fixed effects, and year fixed effects. We pool the data for the two years, restricting the relationship between payment and the explanatory variables to be the same across the two years. Changes in payment shares can be captured by changes in the explanatory variables and by the year fixed effects.

We assume that the relationship between payment shares and explanatory variables is captured by a *fractional multinomial logit* (FM-Logit) model, which states the expected share of each payment type, conditional on the explanatory variables, is a multinomial logit function of the explanatory variables:

$$E[s_{k,j,t} | x_{j,t}] = \frac{\exp(x'_{j,t}\beta_k)}{\sum_{m=1}^4 \exp(x'_{j,t}\beta_m)}, \quad (1)$$

$$k = 1, 2, 3, 4.$$

Before explaining each of the terms in this expression, it will be helpful to understand the subscripts: k and m denote the payment types, cash, debit, credit, and check; j denotes zip code; and t denotes year. The left-hand-side variable, $E[s_{k,j,t} | x_{j,t}]$, is the expected value of the share of type k payments in zip code j in year t , conditional on the time- and location-specific variables $x_{j,t}$ (a vector), which can be thought of as including the state and the year as well as the median transaction size and the demographic and other variables summarized in Table 3. The right-hand side is a function of the explanatory variables as well as coefficients; β_k is a vector of coefficients that multiply the explanatory variables.⁷

By construction, the right-hand side is a number between zero and one as long as the data and coefficients are real numbers. And, by construction, the expected shares always sum to one: $\sum_{k=1}^4 E[s_{k,j,t} | x_{j,t}] = 1$. Note, however, that from (1), for any β_k , $k = 1, 2, 3, 4$, the expected shares are invariant to the transformation $\tilde{\beta}_k = \beta_k + c$, where c is a vector the same length as β_k . In order to achieve identification of β_k , a normalization is needed. We use the standard normalization of setting $\beta_4 = 0$, where $k = 4$ denotes cash. This implies

$$E[s_{4,j,t} | x_{j,t}] = \frac{1}{1 + \sum_{m=1}^3 \exp(x'_{j,t}\beta_m)}. \quad (2)$$

In the Appendix, we present this model in somewhat more detail and explain how the coefficients can be estimated.

⁷ As an alternative to the FMLogit model of payment shares, we could estimate a multinomial logit model at the individual transaction level. By aggregating transactions and modeling shares, we are able to use a larger number of transactions and smooth out the “noise” in individual transactions.

Basic results

We follow the approach described in the Appendix to estimate the model in (1) and (2). In a linear regression model, the usual way to report results is in the form of the estimated coefficients and P-values (or standard errors). With the nonlinear model used here, it is more informative to report *marginal effects* and their P-values; they are presented in Table 4.⁸ For continuous variables, the marginal effect we report (on cash) is the derivative of the predicted share with respect to the variable. For the state and time fixed effects (the former are not reported in the table), the marginal effects we report are the difference between the predicted cash share when the indicator variable is one and when it is zero.

Many of the marginal effects reported in Table 4 are highly significant and have similar magnitudes to those reported in Wang and Wolman (2016a). For example, the median transaction effect is -0.019, compared to -0.018 in the earlier paper. Some of the estimates do differ, however, and not all the marginal effects reported in Table 4 are estimated precisely, in contrast to Wang and Wolman (2016a). The number of different zip codes is roughly comparable in the two papers, but here we use fewer days of data for each zip-code-level observation of the demographic variables. In our earlier paper there were more than 1,000 days of data for each observation of a demographic variable; here there is just one month of data—either February 2011 or February 2015, and this leads to the marginal effects being estimated less precisely.

With respect to age, as discussed above, we interpret the age distribution as the cohort distribution and therefore fix it at its 2011 value. Of course, this means we treat the cohort distribution as fixed so that it cannot explain any of the change in cash shares. In Section 4, we delve into the cohort effect in more detail and present some calculations that represent a rough estimate of the contribution of changes in the cohort distribution to changes in the cash share.

⁸ The dependent variables are the fractions of each of the four general payment instruments used in transactions at stores in a zip code in February 2011 and February 2015. The independent variables take their values in 2011 and 2015. Metro is a dummy variable taking the value of one when the zip code is in an MSA, otherwise it is equal to zero. We rescale some of the variables relative to the levels reported in Table 3 in order to make the marginal effects of common magnitude. Branches per capita is measured as the number of bank branches per 100 residents in a zip code. Robbery rate is defined as the number of robberies per 100 residents in a county. Median household income is measured in units of \$100,000 per household in a zip code. Population density is measured in units of 100,000 residents per square mile in a zip code. All the demographic variables are expressed as fractions.

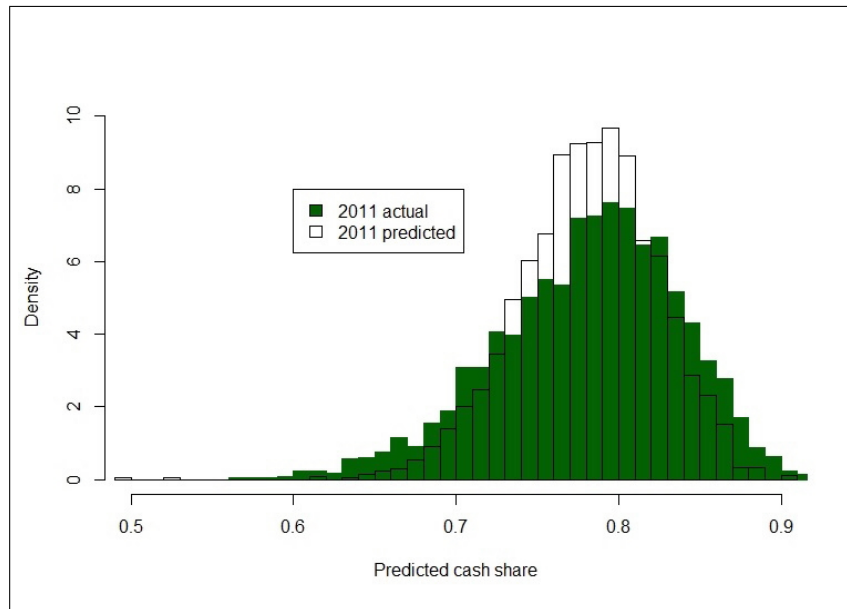
Table 4 Marginal Effects on Cash

Variable	Estimate at mean	P-value
Med. transaction size	-0.019	0.000
(Year=2015) - (Year=2011)	-0.068	0.000
Banking HHI	-0.002	0.469
Banking HHI \times Metro	-0.022	0.000
Branches per capita	-0.040	0.127
Robbery rate	-0.062	0.005
Median household income (\$)	-0.017	0.153
Population density (per mile ²)	0.016	0.535
Family households	-0.089	0.000
Housing: Owner-occupied	-0.364e-04	0.969
Vacant	.013	0.178
Female	-0.027	0.186
Age: 15-34	-0.147	0.000
35-54	-0.114	0.000
55-69	0.016	0.531
≥ 70	6.80e-04	0.981
Race: black	0.063	0.000
Hispanic	0.011	0.050
Native	0.141	0.000
Asian	-0.062	0.007
Pac-Isrl	-0.073	0.627
other	0.009	0.434
multiple	-0.001	0.964
Educ: high school	-0.279	0.000
some college	-0.463	0.000
college	-0.309	0.000

Turning to the model's overall fit, Figures 6 and 7 show that it does a reasonable job of explaining the variation in cash shares across time and locations: Figure 6 compares the actual distribution of 2011 cash shares to the model's predicted distribution, and Figure 7 does the same thing for 2015. The pseudo-R² values are 0.55 for 2011 and 0.59 for 2015.

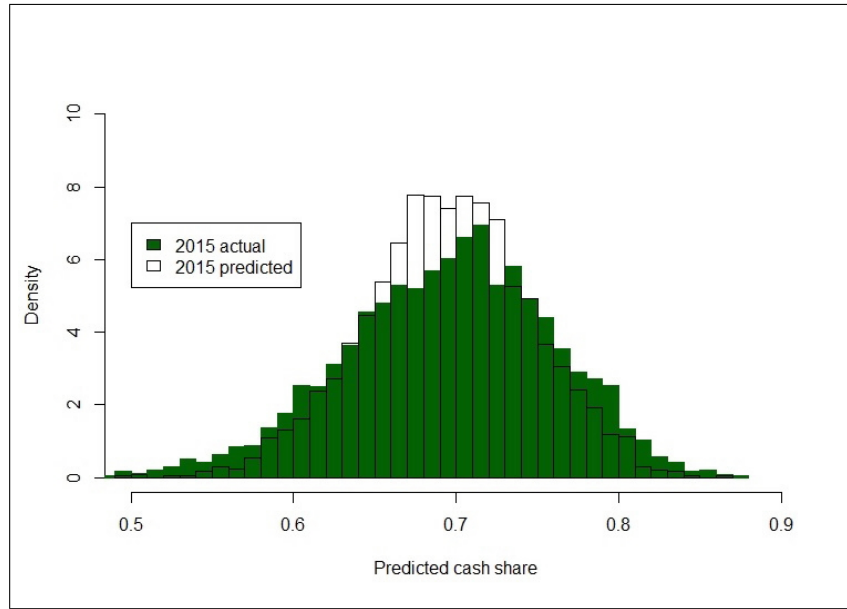
4. ANALYSIS OF DECLINE IN CASH SHARES

Table 1 shows that the mean cash share of transactions declined by 8.7 percentage points from 2011 to 2015. Our model does a good job of capturing this decline: the predicted cash share evaluated at the means of the 2015 data is 8.8 percentage points lower than the predicted cash share evaluated at the means of the 2011 data. Alternatively, we can calculate the predicted cash share for every observation and compare the mean predicted shares for 2011 and 2015: the difference is

Figure 6 2011 Actual (Green) and Predicted Cash Shares

8.7 percentage points. In a linear regression, these two objects would be identical, but because the FMLogit model is nonlinear, the mean predicted value may differ from the predicted value evaluated at the mean of the explanatory variables. We will report both numbers at various points below; they never differ by much.

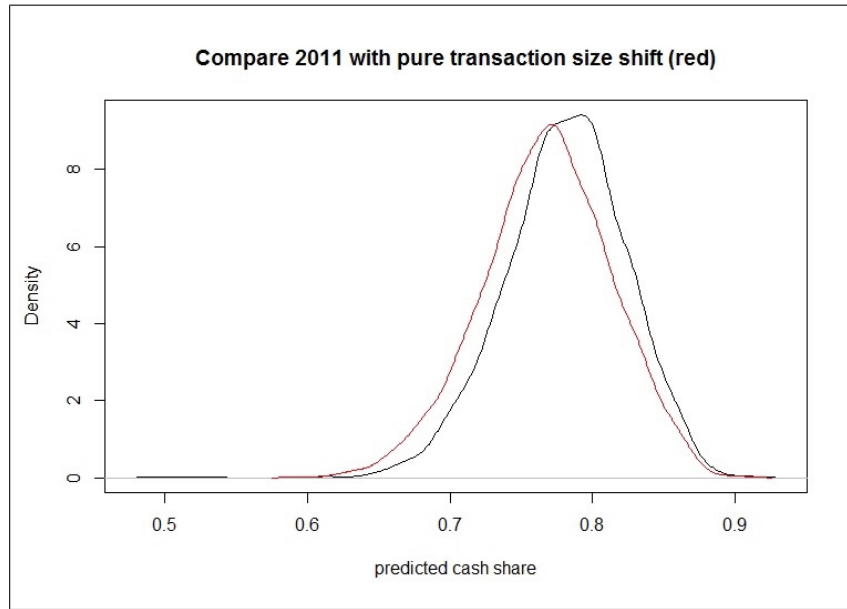
The empirical framework suggests three types of factors to account for the decline in cash shares from 2011 to 2015. First, given a relationship between transaction size and cash shares, an upward shift in the distribution of median transaction sizes (Figure 4) can account for some of the decline in cash shares. Second, given a relationship between demographic variables and cash shares (Table 4), changes in the demographic variables might account for some of the decline in cash shares. And finally, a portion of the decline in cash shares is accounted for by the year dummy; this portion is effectively unexplained and likely attributable to changes in the attributes of noncash payments (e.g., cost, availability, and security) and changing preferences on the part of consumers.

Figure 7 2015 Actual (Green) and Predicted Cash Shares

Increasing average transaction size

The average value of median transaction size increased by \$0.70 from 2011 to 2015. A simple measure of the contribution of changing transaction size to the decline in cash shares is the product of the \$0.70 increase with the marginal effect for transaction size, -0.019 . According to this measure, increasing transaction size can account for a decrease of 1.35 percentage points in the cash share, roughly 15 percent of the total decline. This simple measure ignores nonlinearity in the empirical model. We can take into account the nonlinearity by comparing 2011 predicted cash shares to the shares the model would predict if transaction size changed to its 2015 level but all other variables were fixed at their 2011 values. This approach yields a decrease of 1.33 percentage points in the predicted cash share evaluated at the mean of the explanatory variables and a decrease of 1.44 percentage points in the mean predicted cash share across zip codes. Thus, the linear approximation (1.35 percentage points) turns out to be quite accurate.

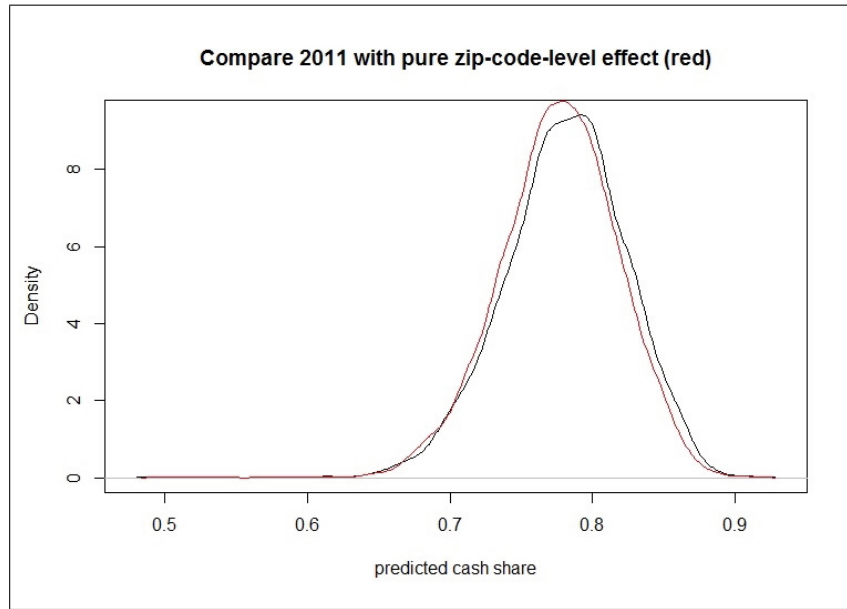
The smoothed density functions in Figure 8 are based on the same approach: the black line represents the density function of predicted cash shares for 2011, whereas the red line represents the density func-

Figure 8 The Transaction Size Effect

tion of counterfactual predicted values, calculated with 2015 transaction size but 2011 values of all other variables. There is a notable leftward shift in the distribution explained by the increase in transaction size, but the shift is small relative to the overall change shown in Figure 2. Note finally that our estimates of the contribution of increasing transaction size to the decrease in cash shares may be affected by correlations between transaction size and some of the zip-code-level variables. This means that a portion of the effect attributed to transaction size could instead be attributed to changes in the zip-code-level variables. In Wang and Wolman (2016b), we explore this idea in more detail by regressing transaction size on the zip-code-level variables and then including the residual portion of transaction size in the FMLogit regression in place of actual transaction size. We find that indeed the marginal effects of other variables change when they include indirect effects of transaction size.

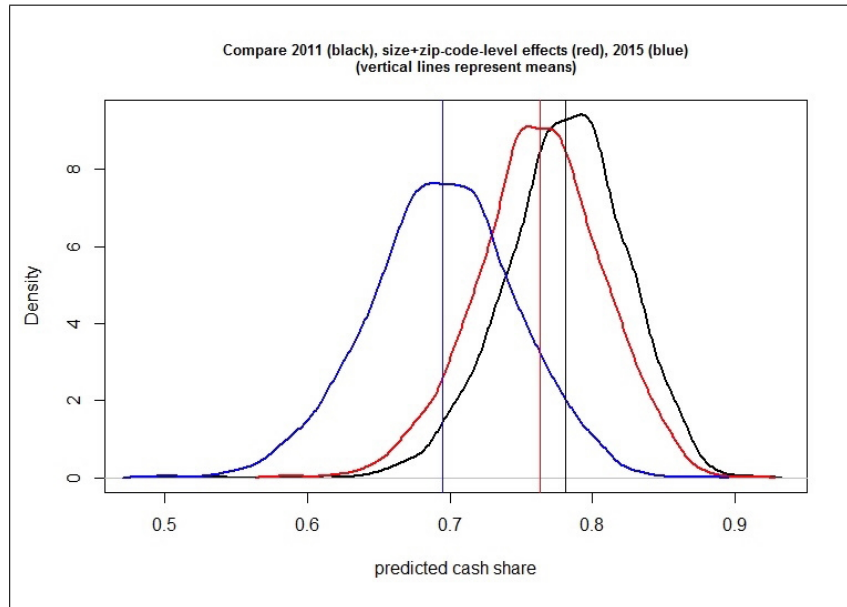
Changing demographic and other variables

Table 4 shows that many location-specific variables have a systematic relationship with the cash share of transactions. Since these variables

Figure 9 The Zip-Code-Level Effect

take on different values in 2011 and 2015, they may be able to account for some of the decline in cash shares over that period. In contrast, Wang and Wolman (2016a) used only a three-year span of data with fixed values of the location-specific variables. As mentioned above, that paper included a rough forecasting exercise that took into account projected changes in the location-specific variables, but the projected changes were identical across locations. In order to quantify the effect of the zip-code-level variables, here we use an analogous approach to that used for transaction size: we compare the predicted cash shares for 2011 with the predicted cash shares implied by holding fixed transaction size and the year dummy at their 2011 values but allowing all the location-specific variables to take on their 2015 values. Comparing the predicted value of cash share conditional on 2011 means to that conditional on 2015 zip-code-level variable means, the 2011 year dummy, and 2011 mean transaction size yields a decline of 0.5 percentage points. This estimate does not change if we instead compare means of predicted values across zip codes.

Figure 9 plots the smoothed density function for 2011 predicted cash shares and compares it to the density of predicted cash shares under the assumption that the zip-code-level variables take on their 2015

Figure 10 Transaction Size and Zip-Code-Level Effects

values but the year dummy and transaction size are fixed at their 2011 values. There is a small but discernible leftward shift in the distribution of predicted cash shares, consistent with the mean estimate. As discussed above, Figures 8 and 9 attribute any effects of transaction size that work through zip-code-level variables to transaction size. In Figure 10, we combine both effects, so that the precise decomposition is irrelevant: the black line is the density of 2011 predicted cash shares; the red line is the density of predicted cash shares holding fixed the year dummy at 2011 but allowing all other variables to change; and the blue line is the density of 2015 predicted cash shares. In Figure 10, the vertical lines represent the respective means. Consistent with our previous calculations, the combination of changes in transaction size and changes in zip-code-level variables accounts for a 1.8 percentage point decline in the mean predicted cash share across zip codes or 1.7 percentage points if we instead use the predicted change in the cash share at the means of the data.

In Wang and Wolman (2016a), the forecasting exercise attributed a relatively large fraction of the projected decrease in the cash share to a cohort effect: a shift in the population toward later-born cohorts who were accustomed to using cards would drive down the cash share of

transactions. Thus far, the calculations here do not take into account that effect because they hold fixed both the age and cohort distribution of the population and the coefficients on age or cohort. Ideally, we would like to treat the cohort distribution just like the other zip-code-level variables in our study: this would involve allowing the cohort distribution to change from 2011 to 2015, estimating a common cohort effect, and then calculating the contribution of the changing cohort distribution to the change in the cash share. The difficulty with this approach is that our data are on age distribution, and in fifteen- and twenty-year bins. Age and cohort are interchangeable at a point in time; for example, the fraction of the population in 2011 that was between 15 and 34 years old (=age) is identical to the fraction of the population in 2011 that was born between 1977 and 1996 (=cohort). However, across time, cohort distributions and age distributions need to be tracked separately unless they are in one-year bins. For example, if we know the fraction of the population that was between 15 and 34 in 2011 and the fraction of the population that was between 15 and 34 in 2015, we have information about two different cohorts in the two years, not the same cohort. For 2011 we have the 1977 to 1996 cohort, and for 2015 we have the 1981 to 2000 cohort. If we knew the age distribution in one-year increments for 2011 and 2015, then it would be trivial to calculate the corresponding cohort distribution in one-year increments.

Without precise data on how the cohort distribution evolved from 2011 to 2015, we nonetheless computed a rough estimate of the contribution of shifts in the cohort distribution to the decrease in cash shares from 2011 to 2015. The idea behind this estimate is to use aggregate census data on a finer gradation of the age distribution to come up with an educated guess about how the cohort distribution changed from 2011 to 2015 across the large bins in our study. Then, we will combine that educated guess with our estimated marginal effects for the different cohorts. Note first that, from Table 4, the cash marginal effect for population aged 35-54 in 2011 is -0.114, compared to 0.016 for age 55-69. The 35-54 age group is the cohort born between 1957 and 1976, and the 55-69 age group is the cohort born between 1942 and 1956. For ages less than 34, the marginal effect is even more negative, and for ages above 69, it is close to zero. According to nationwide census data, the 2011 population share of ages 50-54 was 7 percent. We thus pose the following question: How would the predicted cash share change if there were a 7 percentage point increase in the fraction of the population for whom the cash marginal effect is -0.114, and a 7 percentage-point decrease in the fraction of the population for whom the cash marginal effect is 0.016? The answer is that the predicted cash share would fall by 0.8 percentage points. Adding this to the 1.7

percentage points accounted for by transaction size and other location-specific variables would allow us to account for nearly 30 percent of the overall 8.7 percentage-point predicted decline in the cash share.

The remainder of the predicted decrease in cash shares at the mean of the data—either 7 percentage points or 6.2 percentage points if we include the imputed age effect—is attributed to the year dummy, although this decomposition is not exact: the marginal effect for the year dummy is 6.8 percentage points, and if we compare predicted means for 2011 variables with the year dummy changing, the difference is 6.6 percentage points. Regardless of how we measure it, between 70 and 80 percent of the decline in cash shares cannot be explained by either an increase in transaction size or changes in location-specific variables. We attribute that unexplained decline to a pure “time effect,” which is standing in for all other factors that play a role in payment choice but are not included in the model. The leading candidates for these factors are wider availability, better security, and lower cost of cards, as well as evolving consumer perceptions of each of those factors.

5. CONCLUSION

The cash share of transactions at a large national discount retailer declined by approximately 8.6 percentage points from February 2011 to February 2015. Following up on Wang and Wolman (2016a,b), we use a FMLogit model to study the cash share of transactions across time and locations. The geographic coverage is similar to our earlier paper: thousands of store locations, at the zip-code level. The time coverage is more sparse here: two months, four years apart, as opposed to three years of daily transaction shares in our earlier paper. By restricting the time dimension to low-frequency changes, in this paper we are able to introduce time variation in the zip-code-level variables. Previously, we *measured* the trend decrease in cash shares but were able to attribute it only to a pure time trend or an increase in transaction sizes. We used forecasts of demographic variables to produce a crude measure of the projected contribution of changes in those variables to changes in the cash share. The main contribution of this paper is to explicitly decompose the trend decrease in cash use into a component due to changes in demographic and location-specific variables, as well as a transaction-size component and a pure time effect. We find that location-specific changes in demographic and other variables account for between 0.5 and 1.3 percentage points of the 8.6 percentage-point overall decline. Increasing transaction sizes account for 1.3 percentage points, which leaves between 70 and 80 percent of the decline in cash use unexplained. The unexplained portion is likely being driven

by improved actual characteristics of payment cards as well as slowly evolving consumer perceptions of those characteristics.

Referring back to the introduction, although we attribute a relatively small portion of the decline in cash use to location-specific factors, it would be premature to dismiss distributional arguments about the benefits of currency. First, evaluating those arguments requires quantifying the benefits of currency and payment cards to different groups; that is not part of our analysis and would require an economic model. Second, for the stores and time period in our study, the share of cash transactions declined from 78 percent to 70 percent. Whether our results would carry over to a much larger decline in cash use is an open question, to which time may help provide the answer. Finally, our focus has been on demographic and other location-specific factors across the store locations in our study. As discussed in Wang and Wolman (2016a), those stores are generally located in relatively low-income zip codes. It is possible that analysis of additional retailers in other locations would reveal that demographics account for a greater proportion of the change in cash shares; that is, part of the change in cash shares that we label unexplained may be accounted for by characteristics that are common to the stores and customers studied here but that are distinctive in the context of the entire US economy.

APPENDIX: THE FRACTIONAL MULTINOMIAL LOGIT MODEL

The regression analysis in the paper uses the FMLogit model. The FMLogit model conforms to the multiple fractional nature of the dependent variables, namely that the fraction of payments for each instrument should remain between 0 and 1, and the fractions add up to 1. The FMLogit model is a multivariate generalization of the method proposed by Papke and Wooldridge (1996) for handling univariate fractional response data using quasi-maximum likelihood estimation. Mulahy (2010) provides more econometric details.

Formally, consider a random sample of $i = 1, \dots, N$ zip-code-day observations, each with M outcomes of payment shares. In our context, $M = 4$, which corresponds to cash, debit, credit, and check. Letting s_{ik} represent the k^{th} outcome for observation i , and x_i , $i = 1, \dots, N$, be a vector of exogenous covariates, the nature of our data requires that

$$s_{ik} \in [0, 1] \quad k = 1, \dots, M;$$

$$\Pr(s_{ik} = 0 \mid x_i) \geq 0 \quad \text{and} \quad \Pr(s_{ik} = 1 \mid x_i) \geq 0;$$

$$\text{and} \quad \sum_{m=1}^M s_{im} = 1 \quad \text{for all } i.$$

Given the properties of the data, the FMLogit model provides consistent estimates by enforcing conditions (3) and (4),

$$E[s_k \mid x] = G_k(x; \beta) \in (0, 1), \quad k = 1, \dots, M; \quad (3)$$

$$\sum_{m=1}^M E[s_m \mid x] = 1; \quad (4)$$

and also accommodating conditions (5) and (6),

$$\Pr(s_k = 0 \mid x) \geq 0 \quad k = 1, \dots, M; \quad (5)$$

$$\Pr(s_k = 1 \mid x) \geq 0 \quad k = 1, \dots, M; \quad (6)$$

where $\beta = [\beta_1, \dots, \beta_M]$.⁹ Specifically, the FMLogit model assumes that the M conditional means have a multinomial logit functional form in

⁹ To simplify the notation, we suppress the “ i ” subscript in Eqs (3)-(9).

linear indexes as

$$E[s_k | x] = G_k(x; \beta) = \frac{\exp(x\beta_k)}{\sum_{m=1}^M \exp(x\beta_m)}, \quad k = 1, \dots, M. \quad (7)$$

As with the multinomial logit estimator, one needs to normalize for identification purposes, and we choose the normalization $\beta_M = 0$. Therefore, Eq (7) can be rewritten as

$$G_k(x; \beta) = \frac{\exp(x\beta_k)}{1 + \sum_{m=1}^{M-1} \exp(x\beta_m)}, \quad k = 1, \dots, M - 1; \quad (8)$$

and

$$G_M(x; \beta) = \frac{1}{1 + \sum_{m=1}^{M-1} \exp(x\beta_m)}. \quad (9)$$

Finally, one can define a multinomial logit quasilielihood function $L(\beta)$ that takes the functional forms (8) and (9) and uses the observed shares $s_{ik} \in [0, 1]$ in place of the binary indicator that would otherwise be used by a multinomial logit likelihood function, such that

$$L(\beta) = \prod_{i=1}^N \prod_{m=1}^M G_m(x_i; \beta)^{s_{im}}. \quad (10)$$

The consistency of the resulting parameter estimates $\hat{\beta}$ then follows from the proof in Gourieroux et al. (1984), which ensures a unique maximizer. In our regression analysis, we use Stata code developed by Buis (2008) for estimating the FMLogit model.

REFERENCES

- Buis, Maarten L. 2008. “FMLogit: Stata Module Fitting a Fractional Multinomial Logit Model by Quasi Maximum Likelihood.” Statistical Software Components, Department of Economics, Boston College (June).
- Federal Reserve System. 2013. “The 2013 Federal Reserve Payments Study.” <https://www.federalreserve.gov/paymentsystems/fr-payments-study.htm> (December).
- Federal Reserve System. 2016. “The 2016 Federal Reserve Payments Study.” <https://www.federalreserve.gov/paymentsystems/fr-payments-study.htm> (December).
- Gourieroux, Christian, Alain Monfort, and Alain Trognon. 1984. “Pseudo Maximum Likelihood Methods: Theory.” *Econometrica* 52 (May): 681–700
- Mullahy, John. 2010. “Multivariate Fractional Regression Estimation of Econometric Share Models.” Working Paper 16354. Cambridge, Mass.: National Bureau of Economic Research. (September).
- Papke, Leslie E., and Jeffrey M. Wooldridge. 1996. “Econometric Methods for Fractional Response Variables with an Application to 401(K) Plan Participation Rates.” *Journal of Applied Econometrics* 11 (November/December): 619–32.
- Rogoff, Kenneth S. 2016. *The Curse of Cash: How Large-Denomination Bills Aid Crime and Tax Evasion and Constrain Monetary Policy*. Princeton, N.J.: Princeton University Press.
- Wang, Zhu, and Alexander L. Wolman. 2016a. “Payment Choice and Currency Use: Insights from Two Billion Retail Transactions.” *Journal of Monetary Economics* 84 (December): 94–115.
- Wang, Zhu, and Alexander L. Wolman. 2016b. “Consumer Payment Choice in the Fifth District: Learning from a Retail Chain.” Federal Reserve Bank of Richmond *Economic Quarterly* 102 (First Quarter): 51–78