# Customers and Retail Growth\*

Liran Einav Peter J. Klenow

Jonathan D. Levin Raviv Murciano-Goroff<sup>†</sup>

October 2022

#### Abstract

Using Visa debit and credit card transactions in the U.S. from 2017 to 2021, we document the importance of customers in accounting for sales variation across merchants, across stores within retail chains, and over time for individual merchants and stores. The number of unique customers, as opposed to transactions per customer or dollar sales per transaction, consistently accounts for about 80% of sales variation. The top growing and shrinking merchants account for the majority of total sales reallocation over time, through their acquisition and loss of customers. To illustrate potential implications, we write down an endogenous growth model with a customer margin. In this model, the customer margin drastically increases the size of the largest retailers and their contribution to overall growth, boosting aggregate innovation and growth in the process.

<sup>\*</sup>We are grateful to Emek Basker, Sam Kortum, and Sara Moreira for helpful discussions, and to Jean-Felix Brouillette, Yue Cao, and Mariana Guido for excellent research assistance. Financial support from the National Science Foundation (Grant # SES-1729395) is gratefully acknowledged. Conclusions expressed herein are those of the authors and do not necessarily represent the views of Visa, Inc.

 $<sup>^{\</sup>dagger}\text{Einav},$  Klenow, and Levin: Stanford University and NBER; Murciano-Goroff: Boston University.

# 1. Introduction

Over the last two decades, a stream of research has emphasized the role of customer acquisition in firm dynamics, trade, and growth. Influential models include Fishman and Rob (2003), Luttmer (2006), Arkolakis (2010, 2016), and Perla (2019). Gourio and Rudanko (2014) and Gilchrist, Schoenle, Sim and Za-krajšek (2017) argue that such frictions play a role in business cycle fluctuations.

In this paper, we use Visa debit and credit card transactions from 2017–2021 to bring new systematic and direct evidence to bear on the importance of customers in the U.S. retail sector.<sup>1</sup> The Visa data covers a significant part of consumer spending in the U.S. Roughly 93% of households used at least one debit or credit card in 2018 (Foster, Greene and Stavins, 2019). Around 24% of all U.S. consumer spending flowed through Visa in 2019.<sup>2</sup> Visa accounts for about 60% of all debit and credit card spending, so our analysis might be rerepresentative of 40% of total consumption.<sup>3</sup>

We start by decomposing Visa sales at the chain and store level into the number of unique credit and debit cards, transactions per card, and sales per transaction. We find that the number of customers dominates the decomposition across merchants, across stores within merchants, and over time within stores or merchants. The customer margin is more important for brick-and-mortar transactions than for e-commerce. Focusing on offline retail, we show that about 80% of sales variation can be traced to the number of customers. The importance of customers per store plays an even bigger role than the number of stores for sales variation across merchants and over time for a given merchant.

<sup>&</sup>lt;sup>1</sup>The sample is anonymized. We can link transactions of the same card through a masked card identifier, but the name, street address, or any other confidential personal information about the cardholder is unobserved.

 $<sup>^2</sup>$ Visa (2019)'s 2019 10-K filing reports \$3.242 trillion in nominal payments volume for consumer credit and debit. This is 24.4% of BEA nominal consumption in 2019 of \$13.280 trillion.

<sup>&</sup>lt;sup>3</sup>Consistent with wide spending coverage, in Yelp data for seven mid-sized cities (Pittsburgh, Charlotte, Urbana-Champaign, Phoenix, Las Vegas, Madison, and Cleveland) in 2017, three quarters of the outlets reported payment information and 93% of them indicated that they accepted credit cards (https://www.yelp.com/dataset).

Only for the largest merchants does the store margin (not surprisingly) play a big role. The importance of customers is remarkably consistent across all retail categories, such as furniture, electronics, restaurants, and gas stations.

We continue by showing that around 70% of aggregate sales increases and decreases can be traced to the 1% fastest growing and shrinking merchants (in absolute terms, not percent terms) in a given year. This is consistent with a body of results on the role of fast-growing firms in aggregate job creation, such as Decker, Haltiwanger, Jarmin and Miranda (2016). In the Visa data, we find that most of this tail behavior reflects adding or losing customers.

Our evidence for a large customer margin is in the spirit of Foster, Haltiwanger and Syverson (2008, 2016) and Hottman, Redding and Weinstein (2016). These studies find that fast-growing manufacturers experience rising demand for their products, as opposed to selling a wider array of products more cheaply. One explanation for this could be that such firms are attracting more customers, perhaps linked to the quality of their products.

Eaton, Eslava, Jinkins, Krizan and Tybout (2021) present evidence and a customer search model of exporting firm dynamics. Bornstein (2021) argues that consumer aging interacts with customer inertia to explain the decline in both labor's share and firm entry in recent decades. Baker, Baugh and Sammon (2021) also analyze the customer margin using debit and credit card transactions, from 2010 to 2015. Their focus is on publicly traded firms with observable stock returns. Afrouzi, Drenik and Kim (2021) and Argente, Fitzgerald, Moreira and Priolo (2022) explore the relationship between price-cost markups and margins of consumer demand. Bernard, Dhyne, Magerman, Manova and Moxnes (2022) document the importance of customers for Belgian inter-firm transactions. What sets our study apart from these earlier ones is our focus on retail firms rather than manufacturing firms, as well as the breadth of consumption covered by firms accepting Visa.

To illustrate some of the potential implications of our empirical findings, we write down an endogenous growth model with firm dynamics and a customer

extensive margin. We are modeling productivity growth in the *retail* sector, not the upstream manufacturing sector. This could reflect improvements in store processes and amenities as well as increases in the variety of products offered. Foster, Haltiwanger and Krizan (2002, 2006) trace a nontrivial fraction of aggregate U.S. productivity growth to the Big Box revolution in retail in recent decades. Consistent with this, the U.S. Bureau of Labor Statistics estimates that real labor productivity grew over 3% per year in retail trade from 1988 to 2020.<sup>4</sup>

In our model, retail firms invest in improving their quality, which generates growth in the aggregate. Innovation outcomes are stochastic, so retailers are heterogeneous in their quality levels and growth rates. Retailers spend on marketing to access customers. Because they sell more to each customer they access, retailers with higher quality spend more on marketing and access more customers. Customer acquisition thereby amplifies size differences stemming from quality differences across retailers. Customers are a static function of current year marketing efforts, so retailers have no incentive to lower markups and build their customer base dynamically.<sup>5</sup>

When we calibrate the model to fit the facts we document, the customer margin dramatically amplifies the effect of quality differences and therefore the market share and growth contribution of large retailers. Customers make profits more convex in quality, incentivizing retailers to do more research. This generates faster growth than in an economy in which marketing effort does not rise so quickly with quality and the customer margin is weaker.

The rest of the paper proceeds as follows. Section 2 describes the Visa dataset. Section 3 presents evidence on the importance of customers for sales variation across retailers and their stores. Section 4 lays out the growth model, calibrates it, and describes its quantitative properties. Section 5 concludes.

<sup>&</sup>lt;sup>4</sup>https://www.bls.gov/productivity/tables/labor-productivity-detailed-industries.xlsx

<sup>&</sup>lt;sup>5</sup>This is consistent with evidence in Fitzgerald, Haller and Yedid-Levi (2022) that markups do not rise with age, albeit for Irish exporting firms.

# 2. The Visa dataset

Our primary source of data relies on all credit and debit card transactions that were processed through Visa's electronic payments network in the U.S. between May 2017 and December 2021. The Visa network is the largest network in the market, accounting for about 50% of the credit card transaction volume and about 70% of the debit card volume over this period, with Mastercard, American Express, and Discover sharing the rest.<sup>6</sup>

The unit of observation is a transaction, which includes a merchant identifier, an anonymized card identifier, the time and date of the transaction, and the transaction amount. We do not see the specific items purchased, nor their prices or quantities. The merchant details include an exact store location, so each merchant's store can be uniquely identified.

We apply standard filters used by Visa's data analytics team. We exclude PINdebit transactions (as opposed to signature-debit transactions) because their volume flowing through Visa fluctuates substantially with regulatory changes during our sample period. We also exclude transactions that are not original sales drafts (these would include chargebacks, failed transactions, or payment holds, which would not culminate in an actual transaction), those coming from prepaid gift cards, and those conducted by cards that transacted at fewer than five merchants during the lifetime of the card (these are likely specialized merchantspecific rewards cards). We also exclude transactions associated with merchants located outside the U.S. (which would flow through the U.S. Visa network if the card is issued by a U.S. bank). Online Appendix A provides more detail.

We further restrict the analysis to merchants who are (self) classified as operating in the retail sector (Census Bureau NAICS 44 and 45) or as restaurants (NAICS 722), and we limit our primary analysis to in-person transactions where the card was used in a brick-and-mortar store. Thus, our main sample drops NAICS code 454 ("Nonstore Retail"), which consists almost exclusively of online transactions. We also exclude Gas Stations (NAICS 447) when we decompose

<sup>&</sup>lt;sup>6</sup>https://WalletHub.com/edu/market-share-by-credit-card-network/25531.

aggregate time series changes, given that gasoline sales are heavily driven by price fluctuations (Levin, Lewis and Wolak, 2017).

Overall, the 2017–2021 Visa data contain an annual average of 379 million cards, 30.7 billion transactions, and \$1049 billion in sales for the retail sector plus restaurants.<sup>7</sup> Of these sales, 60% (of the dollar volume) were credit transactions and 40% were debit transactions. Visa spending covers a similar share of retail and restaurant spending in 2021 as in consumption overall. Thus, if other card transactions are similar in nature to Visa's, then Visa spending would be representative of approximately 40% of all retail and restaurant sales.

We analyze the Visa data at three levels of aggregation. First, we aggregate the transaction data to store-card-years to calculate each card's yearly spending in each store. Second, we aggregate to store-years. For every store-year we calculate the following: the number of distinct customer cards (accounts), the number of transactions (swipes), and the dollar volume of transactions. Third, we aggregate to merchant-years. We then calculate, for each merchant, the following variables: number of distinct locations (stores), number of distinct customer accounts (cards), number of transactions, and dollar volume.

Finally, we note that we also have access to Visa data before 2017, going back to 2010, but it is less granular with respect to stores and merchants. For the largest merchants (covering about 70% of the transactions and 60% of the dollar volume during these years), pre-2017 data do not provide exact location for each transaction, only a 5-digit zip code, which makes it infeasible to distinguish stores of the same merchant within a zip. Smaller merchants in these earlier years are grouped by NAICS, so it is also infeasible to distinguish different stores of different small merchants within a NAICS-zip combination, rendering them mostly unusable for our purpose. Therefore, in our main analysis we use the complete set of merchants and stores using data from 2017–2021, but we also report results that use larger merchants only for this longer panel of 2010–2021 (see Online Appendix B).

<sup>&</sup>lt;sup>7</sup>Appendix Table A2 provides these statistics for each year separately. This includes cardpresent transactions only.

# 3. The empirical importance of customers

### **3.1. Sales Decompositions**

**Measurement.** To gauge the importance of customers to a merchant's or store's sales, we decompose sales into three margins we can observe in the Visa data:

$$S = N \cdot \frac{V}{N} \cdot \frac{S}{V},\tag{1}$$

where S denotes total merchant (or store) sales in dollars over a given period, N is the number of unique customer cards that transact at the merchant or store over that period, and V is the total number of visits (transactions) at the merchant or store in that period. The decomposition breaks down total sales into a customer extensive margin (the number of cards) and two intensive margins — the frequency at which customers visit the merchant or store, V/N, and the average transaction amount (the "ticket size"), S/V.

At the merchant level, we can further decompose how merchants reach customers into their number of locations (stores), *L*, and the number of unique customers per store, so that the total decomposition becomes:

$$S = L \cdot \frac{N}{L} \cdot \frac{V}{N} \cdot \frac{S}{V}.$$
 (2)

To operationalize this decomposition, we take logs of both sides in equations (1) or (2) and regress each right-hand-side component on log sales. These coefficients add up to 1 by construction. The coefficients are equivalent to a variance decomposition in which the covariance terms are split equally.

**Overall results.** Table 1 presents this decomposition at the merchant level using different subsamples of merchants in 2019. Panel A reports results from all sectors (that is, not only retail), covering over two million merchants. In this broad sample, the customer margin accounts for 73% of sales variation across merchants, transactions per card account for 4%, and ticket size accounts for 24%. When we look exclusively at online transactions (Panel B), the customer

	Stores	Cards/Store	Trans/Card	Dollar/Trans
A All Data		0 727	0.037	0 236
$(N_{1}, 2, 244, 022)$		(< 0.001)	(< 0.001)	(<0.001)
(1V = 2, 244, 955)		(< 0.001)	(< 0.001)	(<0.001)
		[ 0.734]	[ 0.184]	[ 0.469]
B. Online		0.646	0.049	0.305
(N = 500, 625)		(<0.001)	(<0.001)	(<0.001)
		[ 0.592]	[0.161]	[ 0.430]
C. Offline	0.107	0.642	0.031	0.220
(N = 1, 825, 208)	(<0.001)	(<0.001)	(<0.001)	(<0.001)
	[0.111]	[ 0.694]	[ 0.217]	[ 0.492]
D. Offline Retail	0.106	0.680	0.038	0.176
(N = 928, 678)	(<0.001)	(<0.001)	(<0.001)	(<0.001)
	[ 0.115]	[ 0.745]	[ 0.228]	[ 0.543]
E. Offline Retail, "named"	0.448	0.437	0.024	0.092
(N = 9, 674)	( 0.004)	( 0.004)	( 0.002)	( 0.003)
	[ 0.654]	[ 0.674]	[ 0.284]	[ 0.477]
B. Online (N = 500, 625) C. Offline (N = 1, 825, 208) D. Offline Retail (N = 928, 678) E. Offline Retail, "named" (N = 9, 674)	0.107 (<0.001) [0.111] 0.106 (<0.001) [0.115] 0.448 (0.004) [0.654]	0.646 (<0.001) [0.592] 0.642 (<0.001) [0.694] 0.680 (<0.001) [0.745] 0.437 (0.004) [0.674]	0.049 (<0.001) [0.161] 0.031 (<0.001) [0.217] 0.038 (<0.001) [0.228] 0.024 (0.002) [0.284]	0.305 (<0.001) [0.430] 0.220 (<0.001) [0.492] 0.176 (<0.001) [0.543] 0.092 (0.003) [0.477]

## Table 1: Sales Decomposition for Different Merchant Samples

Notes: N = the number of merchant observations. Cards = the number of unique debit and credit cards; Stores = the number of stores for offline transactions (one for online merchants or at the offline merchant level); Trans = the number of transactions. Standard errors are reported in parentheses, and R-Squared values are in square brackets. Regressions are based on 2019 data. All Data covers all merchants with Visa transactions in consumer NAICS. The "named" merchants are the largest chains. Each regression includes NAICS fixed effects.

margin falls to 65% of variation in online sales across merchants. In contrast, the customer margin accounts for 75% of variation in offline sales across over 1.8 million merchants in 2019 (Panel C). Of this 75%, 64% comes from cards per store and only around 11% from the number of stores.

Our primary focus is on offline retail (plus restaurants), a sector that contains almost a million distinct merchants in 2019. The results (in Panel D) are very similar to those obtained using the broader set of offline merchants. For comparison, the bottom panel (Panel E) shows that for the much smaller set of close to 10,000 large "named" merchants, which Visa tracks all the way back to 2010, the store margin is much more important, accounting for 45% of the variation in sales vs. 44% that is accounted for by cards per store.

In Table 2 we focus on offline retail (plus restaurants), now with more sources of variation. The first row (Panel A) reproduces the cross-sectional analysis we already reported in Panel D of Table 1. The second row (Panel B of Table 2) uses the same set of merchants, over the five years of data (2017–2021), but now focusing on variation in sales over time within each merchant. To do so, we aggregate observation at the merchant-year level (there are over 4.5 million observations at this aggregation level) and include in all regressions merchant and year fixed effects so that the variation is coming from merchants that grow faster or slower than the average for that year. The customer extensive margin is just as important here, accounting for 85% of the variation of sales within merchants. Much of this (69%) is attributed to the changing number of cards per store, and the rest (16%) to store closings and openings.

Panels C and D of Table 2 report a similar analysis at the single store (rather than the merchant) level, where we control for merchant fixed effect in all regressions so that the object of interest is variation in sales across stores within the same merchant. In Panel C we use a cross section of stores (in 2019), and again find that much (84%) of the variation of sales across stores of the same merchant is accounted for by the customer margin. Finally, in Panel D we look at variation in sales within a store over time by (similar to Panel B) using 2017–

	Stores	Cards/Store	Trans/Card	Dollar/Trans
A. Across Merchants	0.106	0.680	0.038	0.176
(N = 928, 678)	[ 0.115]	[ 0.745]	[ 0.228]	[ 0.543]
B. Within Merchants over Time	0.156	0.688	0.103	0.053
(N = 4, 552, 774)	[ 0.751]	[ 0.967]	[ 0.933]	[ 0.970]
C. Across Stores within Merchants		0.837	0.092	0.072
(N = 1, 963, 473)		[ 0.974]	[ 0.814]	[ 0.938]
D. Within Stores over Time		0.815	0.139	0.046
(N = 9, 659, 521)		[ 0.994]	[ 0.964]	[ 0.986]

### Table 2: Decomposing Sales in Offline Retail

Notes: All standard errors are less than 0.001. R-Squared values are reported in square brackets. Across Merchant and Across Store within Merchant decompositions are based on 2019 data. Across Merchant regressions include NAICS fixed effects. Within Merchants over Time and Within Stores over Time are based on 2017–2021 data. Within Merchants over Time regressions include merchant and year fixed effects. Across Stores within Merchants regressions include merchant fixed effects. Within store over Time regressions include store and year fixed effects. See Online Appendix B for robustness with respect to a longer panel of merchant/store data.

2021 data, aggregating variables at the store-year level (we have 9.6 million such observations), and adding store and year fixed effects. The customers margin continues to be the dominant factor (82% in this specification) that explains variation in store sales over time.

Taken together, whether we look across merchants or stores in 2019, or across time for merchants and stores from 2017 to 2021, the number of unique customers explains the vast majority (80% or so) of the variation in sales.

Heterogeneity across retail sectors. In some retail contexts, this general finding seems hardly surprising. For example, in the context of furniture stores, when purchases by a single customer are not frequent, it seems natural that sales are almost entirely driven by how many customers show up. Yet, in other retail contexts this general result is a-priory less obvious. For example, one can imagine that coffee shop sales would be driven not only by how many unique customers show up, but whether they show up once a week or every day, or whether they add a pastry to the coffee.

To explore this, we repeat the decomposition exercise using the "within merchant over time," which is our preferred specification (as in Panel B of Table 2), but estimate it separately for different retail categories (defined by 3-digit NAICS). As before, the observation is at merchant-year level (using data from 2017 to 2021), and each regression includes merchant and year fixed effects.

The results are shown in Figure 1. Customers are the primary driver of merchant sales in all sectors. Customers explain at least 75% of the variation in merchant sales over time in every category. The frequency of visits explains very little of sales variation in these retail categories.

**Households vs. Cards.** Cards could overstate the importance of the customer margin to the extent that households use multiple cards at the same merchant or store — in particular if households use more cards at merchants or stores with higher overall sales. For Visa credit cards from 2016 to 2019 we can match cards to households for about 50% of cards.<sup>8</sup> These households

<sup>&</sup>lt;sup>8</sup>The credit bureau data that allows us to match cards to households was only available to us for transactions between 2016-2019.



Figure 1: Decomposing Merchant Sales Growth by Industry

Notes: This figure displays the coefficients of the "Within Merchant over time" decomposition by industry. The regressions are run with Visa data from 2017 through 2021, and include merchant and year fixed effects.

average 1.7 Visa cards, but 91% of them transact with a given merchant in a given year using a single card. More to the point, when we do a version of Table 2 in Appendix C with cards vs. households, we find the customer margin falls by only one percentage point.

**Non-linearities.** Our linear regressions may hide important non-linearities. We explore this in Figure 2. We partition merchants into 20 bins in terms of their sales (Figure 2a) or sales growth (Figure 2b), with an equal number of merchants in each bin, and plot their components vs. sales (or sales growth) on a log-log base 10 scale. The first bin is normalized to one for all variables. In the cross section of merchants in 2019 (Figure 2a), the number of unique customers is even more important across larger merchants, with visits per cus-

tomer and average transaction amount being less important across the largest merchants. When we look at sales variation over time within a merchant (Figure 2b), after residualizing merchant and year fixed effects, the relationship look approximately linear.



Figure 2: Decomposing Merchant Sales

(a) across Merchants

(b) within Merchants over time

Notes: Panel (a) is based on a cross section of all merchants in 2019. In panel (a), we group the x-axis into 20 bins, and report averages by bin, normalizing each variable by its average for the first bin. Panel (b) repeats the same exercise, but for the panel of merchant-years over 2017–2021. For (b) we de-mean each variable by its merchant average and its year average, so the plot reflects fast vs. slow-growing merchants over time. Both panels are plotted on a log (base 10) scale.

Figure 3 further decomposes the number of customers into the number of stores and the number of cards per store, respectively. Both in the cross section and over time, the number of stores is not an important source of sales variation other than for the largest merchants (the top ventile). This is similar to the role of establishments in firm size more generally, as documented by Moscarini and Postel-Vinay (2012) for example. That is, most variation in firm size comes from its employment per establishment, except for the largest firms which have many more establishments.



Notes: Panel (a) is based on a cross section of all merchants in 2019. In panel (a), we group the x-axis into 20 bins, and report averages by bin, normalizing each variable by its average for the first bin. Panel (b) repeats the same exercise, but uses the panel of merchant-years from 2017–2021. For (b), we de-mean each variable by its merchant average and its year average. Both panels are plotted on a log (base 10) scale.



Figure 4: Decomposing Store Sales

(a) across stores within merchants

(b) within stores over time

Notes: Panel (a) uses a cross section of stores in 2019 and de-means each store by its merchant average. We group the x-axis into 20 bins, and report averages by bin, normalizing each variable by its average for the first group. Panel (b) repeats the same exercise, but uses a panel of stores from 2017–2021, de-meaning each variable by its store average and its year average. All panels are plotted on (base 10) log scale.

## Figure 3: Stores vs. Cards Per Store

Figure 4 repeats this exercise at the store (rather than merchant) level, both for a cross section of stores in 2019 (Figure 4a) and within store over time (Figure 4b). The pattern is remarkably similar for stores and for merchants, except that at the store level the relationships are approximately linear throughout.

## 3.2. Customers and aggregate growth

**Skewed individual contributions to aggregate retail growth.** Having established the importance of the customer margin for growth at the merchant and store levels, we now explore how this translates to retail-wide aggregates.

Let  $S_{it}$  denote merchant *i*'s total sales in year *t*, and  $\Delta S_{it} = S_{i,t} - S_{i,t-1}$  be the change in merchant *i*'s sales from year t - 1 to year *t*. In each year *t*, we order merchants by  $\Delta S_{it}$ , and place them into groups, year by year, which account for the top or bottom 1%, 5%, 10%, or 25% of merchants in terms of their sales change in that year. The top 1% saw the biggest increases in their sales, and the bottom 1% saw the biggest decreases in sales.

We next divide the total increases (or decreases) in each group by the sum of all increases (or decreases) across all merchants in the same year. This is analagous to breaking down the gross job creation and destruction rate as in Davis, Haltiwanger and Schuh (1998), only for the gross sales creation and destruction rates. That is, we trace how much of all sales creation and destruction, respectively, comes from the biggest increases and decreases.

Figure 5 plots the contribution of each group to aggregate sales increases or decreases, averaged across the four observations 2017-2018, ..., 2020-2021. In a similar spirit to Decker et al. (2016), the figure illustrates that a small fraction of growing merchants is responsible for a large fraction of aggregate growth, and similarly a small number of shrinking merchants are responsible for a large fraction of the aggregate decline. The top 1% growers and shrinkers each contribute around 70% of aggregate sales increases and decreases, respectively. The top and bottom 5% contribute more than 80%, the top and bottom 10% contribute



Figure 5: Contribution to Aggregate Sales Changes

Notes: The figure reports the average contribution of each merchant group as defined in the text to aggregate sales change over year with the error bar extending one standard deviation up and down. An observation is a merchant-year and the figure uses a panel of merchants from 2017 to 2021. Each bar corresponds to a merchant group. TX refers to top X% merchants and BX refers to the bottom X% of merchants according to their absolute sales changes.

about 90%, and the top and bottom 25% contribute more than 99%.<sup>9</sup>

In Figure 5 the identity of tail merchants is changing from year to year. How important are *cumulative* sales increases and decreases to the aggregate increases and decreases from 2017 to 2021? To find out we rank merchants based on their cumulative sales changes from 2017 to 2021. This includes entrants among the growers and exiting merchants among the shrinkers.

In Figure 6, we then show that tail merchant contributions remain remarkably similar when looking at cumulative changes from 2017 to 2021. Evidently, many retailers are growing and shrinking by large amounts over the five-year period. This could reflect the short time horizon, but in Online Appendix B we document that the patterns are very similar when we use the longer (2010–2021)

<sup>&</sup>lt;sup>9</sup>In Appendix Figure B10 we show that their initial share of sales is smaller than their share of sales changes. For example, only 45% of sales on average are at the top 1% of growing retailers.



panel, which include a much smaller set of merchants.

Notes: The figure reports the contribution of each merchant group as defined in the text to aggregate sales change between 2017 and 2021. Each bar corresponds to a merchant group. TX refers to the top X% of retailers and BX refers to the bottom X% of retailers by the absolute sales changes between 2017 and 2021.

Figure 7 reports the contributions of the top and bottom 1% of merchants for each 3-digit NAICS from 2017 to 2021. The importance of these tail merchants varies from around 40% in motor vehicles and parts to over 90% for general merchandise, but is mostly in the range of 50% to 80%. Thus, this is a robust feature across retail NAICS that extreme growers and shrinkers account for a large fraction of aggregate sales changes.

The importance of customers for the tails. We now try to assess the extent to which the extensive margin of customers account for these tail patterns. To do so, we decompose merchant sales changes into two components: changes in the number of customers and changes in sales per customer. Let  $N_{it}$  denote the number of unique customers visiting merchant *i* in year *t*, and let  $S_{it}/N_{it}$  denote the sales per customer for merchant *i* in year *t*. Each merchant's sales



Figure 7: Contribution to Aggregate Sales Changes By NAICS

Notes: The figure reports the average contribution of top and bottom 1% merchants to within-NAICs aggregate sales change over year for each retail NAICs. An observation is a merchant-year and the figure uses a panel of merchants from 2017 to 2021. The calculation of each merchant group's contribution is described in the text.

changes can be written as

$$\Delta S_{it} \equiv \Delta N_{it} \cdot \overline{S/N}_{it} + \Delta (S/N)_{it} \cdot \overline{N}_{it}$$

where

$$\overline{N}_{it} \equiv \frac{N_{it} + N_{i,t-1}}{2}$$

and

$$\overline{S/N}_{it} \equiv \frac{S_{it}/N_{it} + S_{i,t-1}/N_{i,t-1}}{2}$$

Using this decomposition, we can attribute the total sales changes in each group to changes in the number of unique customers versus changes in sales per customers. Figure 8 shows that the change in customers accounts for the majority of sales changes in the tails. For decreases in sales, changes in customers accounts for more than 90% of the decline. For increases in sales, changes in customers accounts for just under 60% of the increase.<sup>10</sup> Thus, the (now familiar) pattern prevails even we focus on the tails of the growth/decline distribution: customer growth accounts for most of the extremes we see in overall sales growth across merchants from year to year.



Figure 8: Customers vs. sales/customer and merchant sales changes
Offline Retail, 2018-2021

Notes: The figure reports the average share of sales changes in each merchant group that can be attributed to changes in number of customers and changes in sales per customer respectively from 2018 to 2021. By construction, the two shares sum to 1. Each bar corresponds to a merchant group. TX refers to the top X% of retailers and BX refers to bottom X% of retailers by absolute sales changes.

**Customer turnover.** Given the importance of customers for a merchant's sales and sales growth, a natural question is whether customers shop persistently at the same retailers over time. Figure 9 shows the percent of sales at retailers associated with cards that shopped at the same retailer in the prior

<sup>&</sup>lt;sup>10</sup>The asymmetry between those firms that increased in sales and those that decreased sales is less pronounced in the years prior to the COVID-19 pandemic. In Figure E1, we reproduce this figure using data from the years 2017–2019 and find that changes in customers accounts for approximately 80% of both sales increases and decreases.

year. For this analysis we include only cards that are active in both years. Across all industries, the fraction is 60%, which means that 40% of sales are associated with new customers (relative to the previous year). Thus, there is a fair bit of customer turnover from year to year, even weighted by customer purchases. These findings influence our model, described the in the next section, to model customer acquisition as a static rather than dynamic problem.



Figure 9: Sales by Returning Customers as a percentage of total merchant sales

Notes: The figure reports the average share of sales at merchants in the years 2018-2021 that can be attributed to customers who had shopped at the same merchant in the previous year. We construct this percentage by dividing the total sales of returning customers with the total sales for retailer in each year and then taking the average across years. Each bar corresponds to merchants in a NAICS group.

# 4. A model of growth with customer acquisition

Having shown that the customer margin is quantitatively important, we present a model of growth that incorporates this margin to see how it may matter. The model is purely illustrative; it is certainly possible to write alternative models of marketing and innovation, which would lead to different results.

### 4.1. Customers

Consider a unit mass of customers with identical preferences:

$$U = \sum_{t=0}^{\infty} \beta^{t} \frac{C_{t}^{1-1/\sigma}}{1 - 1/\sigma}.$$

where  $\sigma > 0$  is the intertemporal elasticity of substitution and  $0 < \beta < 1$  is the discount factor. Composite consumption *C* is a CES aggregate of varieties:

$$C_t = \left(\int_0^1 n_{it} \left(q_{it}c_{it}\right)^{\frac{\theta-1}{\theta}} \mathbf{d}i\right)^{\frac{\theta}{\theta-1}}$$

where  $n_{it} \in [0, 1]$  is the probability that a customer purchases from retailer *i* and  $q_{it}$  is the quality of retailer *i*.  $\theta > 1$  is the elasticity of substitution between retailers. Note there is a fixed unit measure of retailers. We assume that  $n_{it}$  is identical across consumers, so it is also the fraction of consumers who buy from retailer *i* in period *t*. Demand (per customer) conditional on access to retailer *i* is given by:

$$c_{it} = \left(\frac{P_t}{p_{it}}\right)^{\theta} q_{it}^{\theta-1} C_t, \quad \forall i \in [0, 1],$$

where the ideal consumer price index is:

$$P_t \equiv \left(\int_0^1 n_{it} \left(\frac{p_{it}}{q_{it}}\right)^{1-\theta} \mathbf{d}i\right)^{\frac{1}{1-\theta}}.$$

Total quantity demanded from retailer *i*, summed across customers, is thus:

$$y_{it} = n_{it}c_{it}.$$

## 4.2. Retailers

Each retailers uses "production" labor  $l_{it}$  to sell to its customers:

$$y_{it} = l_{it}.$$

It uses marketing labor  $m_{it}$  to reach a random fraction  $n_{it}$  of customers:

$$n_{it} = \left(\frac{\gamma m_{it}}{\phi}\right)^{\frac{1}{\gamma}}.$$

Here  $\gamma > 1$  governs the convexity customer acquisition costs, and  $\phi > 0$  the level of marketing costs. We assume no negative externality with respect to the marketing efforts of other retailers.<sup>11</sup>

Choosing labor as the numeraire, retailer static profit maximization is:

$$\max_{p_{it},m_{it}} \left( p_{it} - 1 \right) y_{it} - m_{it}.$$

Assuming that retailers engage in monopolistic competition, they set their price to a constant markup above unit marginal cost:

$$p_{it} = \mu$$
 where  $\mu \equiv \frac{\theta}{\theta - 1}$ .

Substituting the retailer's price in its demand function yields:

$$c_{it} = \left(\frac{q_{it}P_t}{\mu}\right)^{\theta-1} \cdot \frac{P_tC_t}{\mu}.$$

<sup>&</sup>lt;sup>11</sup>Alternatively, we could specify a negative externality as  $n_{it} = \left(\frac{\gamma m_{it}}{\phi M_t^{\delta}}\right)^{\frac{1}{\gamma}}$  where  $M_t \equiv \int_0^1 m_{it} di$  is aggregate marketing labor across all retailers and  $\delta > 0$  controls the magnitude of the negative externality. Doing so would not affect the quantitative results we obtain from the model below. Online Appendix F contains the model with a parameterized marketing externality.

The retailer's static marketing problem becomes:

$$\max_{n_{it}} n_{it} \left(\frac{q_{it}P_t}{\mu}\right)^{\theta-1} \cdot \frac{P_t C_t}{\theta} - \frac{\phi n_{it}^{\gamma}}{\gamma}.$$

This marketing problem yields the following first order condition:

$$n_{it} = \min\left\{ \left(\frac{q_{it}P_t}{\mu}\right)^{\theta-1} \cdot \frac{P_t C_t}{\theta \phi}, 1 \right\}^{\frac{1}{\gamma-1}}.$$

Denoting  $\Gamma \equiv \frac{\gamma}{\gamma-1}$ , it follows that a retailer's flow profits are:

$$\pi_{it} = \left[ \left( \frac{q_{it}P_t}{\mu} \right)^{\theta-1} \cdot \frac{P_t C_t}{\theta} \right]^{\Gamma} \cdot \frac{\phi}{\Gamma}.$$

It is useful to define an aggregate quality index as:

$$Q_t \equiv \left(\int_0^1 q_{it}^{\Gamma(\theta-1)} \mathbf{d}i\right)^{rac{1}{\Gamma(\theta-1)}}.$$

Defining a retailer's relative quality as  $z_{it} = q_{it}/Q_t$ , we can use the market clearing conditions for the final good and labor to solve for profits as a function of aggregate production labor:

$$\pi_{it} = \frac{L_t z_{it}^{\Gamma(\theta-1)}}{\Gamma(\theta-1)}.$$

## 4.3. Innovation

A retailer with absolute quality  $q_{it}$  and relative quality  $z_{it}$  that hires research labor  $s_{it}$  sees its quality follow a controlled binomial process with probability  $x_{it} \in [0, 1]$ :

$$q_{it+1} = \begin{cases} q_{it}e^{\Delta} & \text{with probability } x_{it} \\ q_{it} & \text{with probability } 1 - x_{it} \end{cases} \text{ and } s_{it} = \lambda \log\left(\frac{1}{1 - x_{it}}\right) z_{it}^{\zeta}.$$

Here  $\Delta$ ,  $\lambda$  and  $\zeta$  are all strictly positive.  $\Delta$  is the percentage step size of quality innovations, and  $x_{it}$  is the probability that a retailer succeeds in innovating.  $\lambda$  scales the level of research labor and  $\zeta$  quantifies how much more research labor is necessary to innovate from a higher level of relative quality. Note the knowledge spillover implicit in this formulation: the higher the quality of other retailers, the less labor required to successfully innovate ( $\zeta > 0$ ). A retailer's value function is given by:

$$V_{t}(z) = \pi_{t}(z) + \max_{x \in [0,1]} \left\{ R_{t}^{-1} \left[ x V_{t+1} \left( z e^{\Delta - g_{t}} \right) + (1-x) V_{t+1} \left( z e^{-g_{t}} \right) \right] - s(z,x) \right\}$$

where *R* is the gross interest rate and *g* is the growth rate of aggregate quality:

$$1 + g_t = \left(1 + \int x(z) z^{\Gamma(\theta-1)} \left(e^{\Delta \Gamma(\theta-1)} - 1\right) \mathbf{d} F_t(z)\right)^{\frac{1}{\Gamma(\theta-1)}}.$$

The Euler equation produces the usual relationship between the growth rate g and the consumer's discount factor in the absence of aggregate uncertainty:

$$\left(1+g_t\right)^{1/\sigma} = \beta R_t.$$

Note that the growth rate of consumption g is the same as that for the aggregate quality index in a stationary equilibrium. Meanwhile, the first-order condition of the retailer's dynamic research problem implies:

$$x_t(z) = 1 - \frac{\lambda z^{\zeta} R_t}{V_{t+1}(ze^{\Delta - g_t}) + V_{t+1}(ze^{-g_t})}.$$

## 4.4. Labor market clearing

To recap, labor is used for production, marketing, and research:

$$L_{t} = \int l_{t}(z) \, \mathrm{d}F_{t}(z)$$
$$M_{t} = \int m_{t}(z) \, \mathrm{d}F_{t}(z)$$
$$S_{t} = \int s_{t}(z) \, \mathrm{d}F_{t}(z).$$

As each of the unit mass of consumers is endowed with one unit of labor that they supply inelastically, the labor market clearing condition is simply:

$$L_t + M_t + S_t = 1.$$

Solving for aggregate production and marketing labor yields:

$$L_t = rac{\gamma \left( heta - 1 
ight) \left( 1 - S_t 
ight)}{\gamma \left( heta - 1 
ight) + 1} \quad ext{and} \quad M_t = rac{L_t}{\gamma \left( heta - 1 
ight)}.$$

## 4.5. Calibration

A period in the model is one year. Without loss of generality we set the aggregate labor endowment to  $1.^{12}$  In Table 3 we present our baseline parameter values, which we choose as follows:

- We set the intertemporal elasticity of substitution  $\sigma$  = 0.5 based on Hall (2009).
- We choose an elasticity of substitution between varieties  $\theta$  to 3. This is at the lower end of estimates such as in Hottman et al. (2016), but this and other papers typically do not control for the customer margin.
- We set the discount factor  $\beta$  to 0.992 so that, when the baseline growth rate is set to 2.9% per year to match the data (see below), the steady state

<sup>&</sup>lt;sup>12</sup>We set research cost parameters to match the target growth rate given the labor endowment.

real interest rate is 6.7% per year as in Farhi and Gourio (2018).

- We set the level of marketing costs φ so that the retailer with maximum relative quality reaches half of the customers.<sup>13</sup> The resulting cost of increasing the share of customers reached (n) by one percentage point is equal to 0.34% of revenue for the average retailer.
- We set the elasticity of marketing labor with respect to customers to  $\gamma = 1.25$ . The elasticity of sales with respect to quality in the model is the sum of the elasticity of customers and elasticity of spending per customer with respect to quality:

$$\xi_{y,q} = \xi_{n,q} + \xi_{c,q} = \frac{\theta - 1}{\gamma - 1} + \theta - 1.$$

With  $\gamma = 1.25$  and  $\theta = 3$ , the customer share of the sales elasticity is 80%, which matches our finding in Section 3.

- We choose an innovation step size Δ of 6%, roughly double the growth rate. As a consequence the average probability of innovation success (*x*) is around one half. With θ = 3, γ = 1.25, and Δ = 0.06, sales grow by 29% (g · ξ<sub>y,q</sub>) for expanding retailers and shrink by 29% for contracting retailers.
- We choose the innovation cost parameters  $\lambda$  (the scale of R&D costs) and  $\zeta$  (the research spillover or convexity parameter) to achieve a 3.05% growth rate and for the top 1% fastest growing retailers to account for 70% of sales changes.<sup>14</sup> The 70% figure is from the Visa data, as described above. With these parameter values, the resulting cost of increasing the probability of research success (*x*) by one percentage point is equal to 0.16% of revenue for the average retailer in the model.

<sup>&</sup>lt;sup>13</sup>For confidentiality reasons we cannot target the precise share of customers reached by top merchants in the Visa data.

<sup>&</sup>lt;sup>14</sup>The growth rate is for labor productivity in the retail sector from 1988–2020 according to the U.S. Bureau of Labor Statistics. https://www.bls.gov/productivity/tables/ labor-productivity-detailed-industries.xlsx

Symbol	Parameter	Value
σ	Intertemporal elasticity of substitution	0.50
heta	Elasticity of substitution between varieties	3.00
$\beta$	Discount factor	0.992
$\phi$	Scale of marketing costs	$1.83\cdot 10^{35}$
$\gamma$	Elasticity of marketing costs wrt customers	1.25
$\Delta$	Quality step size	0.063
$\lambda$	Linear research cost	0.094
ζ	Research spillover parameter	10.04

Table 3: Parameter Values

### 4.6. Results

We are now ready to characterize equilibrium outcomes. For contrast we will show what happens in an economy in which the customer margin is much less responsive to a retailer's quality. We achieve this by subsidizing the marketing of low quality retailers and taxing the marketing of high quality retailers so that aggregate marketing labor is the same as in the baseline economy. We keep all other parameter values the same when we make this comparison.

That is, we make the retailer's marketing costs  $\tau m^{\omega} \cdot m$  rather than simply m. Here  $\tau \gg 1$  and  $\omega > 0$ . I.e., we implement a progressive marketing tax. We set  $\omega = 3$  so that the elasticity of customers with respect to sales is only 20% rather than the 80% in the baseline. And we set  $\tau$  so that total marketing labor is unaffected by the marketing subsidy/tax. The tax is not revenue neutral, but we assess a a lump-sum tax on households.

Figure 10a shows how the fraction of consumers the retailer sells to (*n*) varies with the retailer's relative quality (*z*). It is log-linear with elasticity  $\Gamma$  in the baseline and  $\Gamma/(1 + \omega)$  under the progressive marketing tax. The marketing tax thus makes profits less convex in a retailer's relative quality — see Figure 10b.



Figure 10: Customers and retailer value

Notes: In panels (a) and (b), this figure shows how the fraction n of customers reached by the retailer and its value v, vary with it's relative quality z, respectively. In the marketing tax economy marketing costs are  $\tau m^{\omega} \cdot m$  with  $\tau >> 1$  and  $\omega = 3$ , whereas they are m in the baseline economy.





Notes: This figure shows how the arrival rate of innovations x varies with the retailer's relative quality z. In the marketing tax economy marketing costs are  $\tau m^{\omega} \cdot m$  with  $\tau >> 1$  and  $\omega = 3$ , whereas they are m in the baseline economy.

Because customers makes variable profits increase rapidly in quality, higher quality retailers to do more innovation than otherwise (see Figure 11a). A corollary is that R&D intensity (research spending as a share of sales) decreases modestly with z in the heart of the distribution in the baseline case. In contrast, research effort falls sharply with z with a progressive marketing tax. This shifts the bulk of R&D away from large retailers, as depicted in Figure 11b.

Because higher quality retailers do more R&D in the baseline, the stationary distribution of relative qualities is much more dispersed in our baseline than with a marketing tax (Figure 12a).<sup>15</sup> In turn, the distribution of sales across merchants is dramatically more dispersed when there is a strong extensive margin for customers as in the baseline than with a marketing tax (Figure 12b).



#### Figure 12: Distribution of quality and sales

Notes: This figure shows the density of merchant relative quality z in panel (a) and of merchant sales y in panel (b). In the marketing tax economy marketing costs are  $\tau m^{\omega} \cdot m$  with  $\tau >> 1$  and  $\omega = 3$ , whereas they are m in the baseline economy.

In Table 4 we compare the steady state values of variables in the two economies. The growth rate of aggregate quality falls from 3.05% in the baseline to 2.79% with a progressive marketing tax. The real interest rate falls with the growth rate. Marketing labor stays the same by construction. Research labor edges up by 5 basis points at the expense of production labor.

<sup>&</sup>lt;sup>15</sup>In both models the probability of successfully innovating is equal to one for the smallest retailers and zero for the largest ones, which delivers a stationary distribution of relative quality.

Symbol	Variable	Baseline	Marketing tax
g	Growth rate	3.05%	2.79%
r	Interest rate	7.01%	6.47%
L	Production labor	67.9%	67.9%
M	Marketing labor	27.2%	27.2%
S	Research labor	4.92%	4.97%

Table 4: Steady-state endogenous variables

Notes: The baseline economy features the parameter values in Table 3. In the marketing tax economy marketing costs are  $\tau m^{\omega} \cdot m$  with  $\tau >> 1$  and  $\omega = 3$ , whereas they are *m* in the baseline economy.

Why does the marketing tax economy grow more slowly than the baseline economy despite devoting a little more effort to research? The answer lies in the benefits of spreading quality improvements over more customers. To illustrate this we calculate a Törnqvist approximation to the true growth rate:

$$g_t \approx \int \left[\overline{\omega}_t(z) \times \frac{\Delta q_{t+1}(z)}{q_t(z)}\right] \mathbf{d}F_t(z) \quad \text{where} \quad \overline{\omega}_t(z) \equiv \frac{\omega_{t+1}(z) + \omega_t(z)}{2}.$$

This is akin to using a Törnqvist index of input growth rates, only here it is an index of retailer quality growth rates.<sup>16</sup> The Törnqvist weights are the retailer's average sales share across the two years, denoted  $\overline{\omega}_t(z)$ . This approximation to the true growth rate can be decomposed into 1st order and 2nd order terms:

$$g_t \approx \underbrace{\int \left[\omega_t(z) \times \frac{\Delta q_{t+1}(z)}{q_t(z)}\right] \mathrm{d}F_t(z)}_{\text{1st order term}} + \underbrace{\int \left[\left(\overline{\omega}_t(z) - \omega_t(z)\right) \times \frac{\Delta q_{t+1}(z)}{q_t(z)}\right] \mathrm{d}F_t(z)}_{\text{2nd order term}}.$$

Using these expressions, Table 5 shows the importance of the covariance between sales shares and quality growth to aggregate growth. The Table indicates that the Törnqvist second-order approximation is quite good. If we instead use

<sup>&</sup>lt;sup>16</sup>For an example see https://www.bls.gov/opub/hom/inp/calculation.htm.

initial sales shares as weights (i.e., use just the 1st order term, a 1st order approximation), the approximated growth rate understates the true growth rate by half a percentage point in the baseline economy. With the progressive marketing tax, in contrast, the second order term is much smaller. The first order term is higher with the marketing tax than with no marketing tax (2.72% versus 2.60%). So the second order term is entirely responsible for the lower growth rate with a marketing tax. This breakdown illustrates the importance of customer acquisition in amplifying the growth contribution of the right tail of merchant growers. Retailer size dispersion is much higher with customer acquisition, so right tail growers contribute much more to growth in the baseline economy.

	Baseline	Marketing tax
True growth rate	3.05%	2.79%
Approximated growth rate	3.10%	2.84%
1st order term	2.60%	2.72%
2nd order term	0.50%	0.12%

Table 5: Törnqvist growth decomposition

Notes: Each bar corresponds to a merchant group. TX refers to top X% merchants and BX refers to the bottom X% of merchants according to their absolute sales changes. The baseline economy features the parameter values in Table 3. In the marketing tax economy marketing costs are  $\tau m^{\omega} \cdot m$  with  $\tau >> 1$  and  $\omega = 3$ , whereas they are *m* in the baseline economy.

We can also calculate the contribution of the top 1% of growing retailers to aggregate sales increases. As depicted in Figure 13, our baseline model is calibrated to achieve a contribution of 70% from the top 1%, as seen in the Visa data. With a progressive marketing tax and a weak customer margin, in contrast, the top 1% of retailers account for less than 1% of all sales increases. Again, this comes from both the direct effect of acquiring customers in response to rising *z*, and the indirect effect of a much narrower *z* distribution in the absence of a customer margin. In this economy, the growth contribution of top retailers is the same as their contribution to aggregate sales increases. Therefore, the top 1% fastest growing retailers also account for 70% of consumption growth in the baseline economy and less than 1% in the marketing tax economy.



Figure 13: Retailer Contributions to Aggregate Sales Changes

Notes: The baseline economy feature the parameter values in Table 3. In the marketing tax economy marketing costs are  $\tau m^{\omega} \cdot m$  with  $\tau >> 1$  and  $\omega = 3$ , whereas they are *m* in the baseline economy.

Finally, for a given distribution of qualities across retailers, we can decompose the difference in aggregate consumption between the baseline economy and the economy with a progressive marketing tax into two terms, the effective variety of retailers accessed and the quantity bought per retailer:

$$\left[\frac{\int_{0}^{1} n_{it}(q_{it}c_{it})^{\frac{\theta-1}{\theta}}\mathbf{d}i}{\int_{0}^{1} \tilde{n}_{it}(q_{it}\tilde{c}_{it})^{\frac{\theta-1}{\theta}}\mathbf{d}i}\right]^{\frac{\theta}{\theta-1}} = \underbrace{\left[\frac{\int_{0}^{1} n_{it}(q_{it}c_{it})^{\frac{\theta-1}{\theta}}\mathbf{d}i}{\int_{0}^{1} \tilde{n}_{it}(q_{it}c_{it})^{\frac{\theta-1}{\theta}}\mathbf{d}i}\right]^{\frac{\theta}{\theta-1}}}_{\text{Effective variety}} \times \underbrace{\left[\frac{\int_{0}^{1} \tilde{n}_{it}(q_{it}c_{it})^{\frac{\theta-1}{\theta}}\mathbf{d}i}{\int_{0}^{1} \tilde{n}_{it}(q_{it}\tilde{c}_{it})^{\frac{\theta-1}{\theta}}\mathbf{d}i}\right]^{\frac{\theta}{\theta-1}}}_{\text{Ouantity}}.$$

We evaluate these expressions at quality distribution in the baseline economy. Variables without tildes are baseline values, and those with tildes are values under a marketing tax. The effective variety term is equal to 2.92, so that consumers benefit amply from accessing more retailers in the baseline economy. The quantity term is term is equal to 1.002, so there is essentially no difference in quantities.<sup>17</sup> Thus the customer margin provides static variety benefits along with dynamic quality growth benefits in our model.

## 5. Conclusion

Using Visa data on credit and debit card transactions at U.S. retail merchants from 2017 to 2021, we documented the paramount importance of the extensive customer margin in driving variation in retail sales. Customers account for approximately 80% of the sales variation whether we look across merchants, across stores within merchants, or over time within merchants and stores.

We wrote down a simple growth model with a customer extensive margin. In the model, retailers pay marketing and research costs to acquire customers and improve their quality. The customer margin makes large retailers drastically more important for sales and overall growth. The ability to increase profits by adding customers increases the returns to innovation and stimulates growth.

Our model illustrates one of many potential implications of a paramount customer margin. We assumed marketing and R&D expenditures are complements, but see Cavenaile and Roldan-Blanco (2021) for a theoretical and empirical analysis of complementarity versus substitutability between them. More generally, the magnitude of the customer margin is likely to have important implications for the nature and intensity of competition. Loyal customers may increase markups and exacerbate market power. Moreover, the dynamic nature of customer relationships (which we abstract from in the current paper) may increase barriers to entry and in turn slow down the exit of inefficient firms. Increased access to datasets like Visa could allow a more direct exploration of these forces, which we view as promising avenues for future work.

<sup>&</sup>lt;sup>17</sup>The marketing tax is distortionary here. Though that could obviously change if retailers' marketing efforts canceled each other out sufficiently.

# References

- Afrouzi, Hassan, Andres Drenik, and Ryan Kim, "Growing by the masses: Revisiting the link between firm size and market power," 2021.
- Argente, David, Doireann Fitzgerald, Sara Moreira, and Anthony Priolo, "How do entrants build market share? The role of demand frictions," 2022.
- Arkolakis, Costas, "Market penetration costs and the new consumers margin in international trade," *Journal of Political Economy*, 2010, *118* (6), 1151–1199.
- \_ , "A unified theory of firm selection and growth," *The Quarterly Journal of Economics*, 2016, *131* (1), 89–155.
- Baker, Scott R, Brian Baugh, and Marco C Sammon, "Measuring customer churn and interconnectedness," 2021. NBER working paper 27707.
- Bernard, Andrew B, Emmanuel Dhyne, Glenn Magerman, Kalina Manova, and Andreas Moxnes, "The origins of firm heterogeneity: A production network approach," *Journal of Political Economy*, 2022, *130* (7), 1768–1804.
- Bornstein, Gideon, "Entry and profits in an aging economy: The role of consumer inertia," 2021.
- Cavenaile, Laurent and Pau Roldan-Blanco, "Advertising, innovation, and economic growth," *American Economic Journal: Macroeconomics*, 2021, *13* (3), 251–303.
- Davis, Steven J, John C Haltiwanger, and Scott Schuh, "Job creation and destruction," *MIT Press Books*, 1998.
- Decker, Ryan A, John Haltiwanger, Ron S Jarmin, and Javier Miranda, "Where has all the skewness gone? The decline in high-growth (young) firms in the US," *European Economic Review*, 2016, *86*, 4–23.

- Eaton, Jonathan, Marcela Eslava, David Jinkins, Cornell J Krizan, and James R Tybout, "A search and learning model of export dynamics," 2021. NBER working paper 29100.
- Farhi, Emmanuel and Francois Gourio, "Accounting for macro-finance trends," Brookings Papers on Economic Activity, 2018, 2, 147–223.
- Fishman, Arthur and Rafael Rob, "Consumer inertia, firm growth and industry dynamics," *Journal of Economic theory*, 2003, *109* (1), 24–38.
- Fitzgerald, Doireann, Stefanie Haller, and Yaniv Yedid-Levi, "How exporters grow," 2022. NBER Working Paper 21935.
- Foster, Kevin, Claire Greene, and Joanna Stavins, "The 2018 survey of consumer payment choice: Summary results," *Federal Reserve Bank of Atlanta, Research Data Reports*, 2019, (19-02).
- Foster, Lucia, John C Haltiwanger, and Cornell J Krizan, "The link between aggregate and micro productivity growth: Evidence from retail trade," 2002. NBER working paper 9120.
- \_ , John Haltiwanger, and Chad Syverson, "Reallocation, firm turnover, and efficiency: Selection on productivity or profitability?," *American Economic Review*, 2008, 98 (1), 394–425.
- \_ , \_ , and \_ , "The slow growth of new plants: Learning about demand?," *Economica*, 2016, *83* (329), 91–129.
- \_\_\_\_\_, \_\_\_\_, and Cornell J Krizan, "Market selection, reallocation, and restructuring in the US retail trade sector in the 1990s," *The Review of Economics and Statistics*, 2006, *88* (4), 748–758.
- Gilchrist, Simon, Raphael Schoenle, Jae Sim, and Egon Zakrajšek, "Inflation dynamics during the financial crisis," *American Economic Review*, 2017, *107* (3), 785–823.

- Gourio, Francois and Leena Rudanko, "Customer capital," *Review of Economic Studies*, 2014, *81* (3), 1102–1136.
- Hall, Robert E, "Reconciling cyclical movements in the marginal value of time and the marginal product of labor," *Journal of Political Economy*, 2009, *117* (2), 281–323.
- Hottman, Colin J, Stephen J Redding, and David E Weinstein, "Quantifying the sources of firm heterogeneity," *The Quarterly Journal of Economics*, 2016, *131* (3), 1291–1364.
- Levin, Laurence, Matthew S Lewis, and Frank A Wolak, "High frequency evidence on the demand for gasoline," *American Economic Journal: Economic Policy*, 2017, 9 (3), 314–47.
- Luttmer, Erzo G.J., "Consumer search and firm growth," 2006. Working paper, University of Minnesota.
- Moscarini, Giuseppe and Fabien Postel-Vinay, "The contribution of large and small employers to job creation in times of high and low unemployment," *American Economic Review*, 2012, *102* (6), 2509–2539.
- Perla, Jesse, "A model of product awareness and industry life cycles," 2019. Working paper, University of British Columbia.
- Visa, Inc., "Visa, Inc., 2019 Annual Report, 10-Filing," https://investor.visa. com/SEC-Filings/ November 2019.