

# Customers and Retail Growth\*

Liran Einav Peter J. Klenow

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

December 2021

## Abstract

Using Visa debit and credit card transactions in the U.S. from 2016 to 2019, 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. Customers, as opposed to transactions per customer or dollar sales per transaction, consistently account for about 80% of sales variation. The top 1% of growing and shrinking merchants account for about 70% of customer and sales reallocation in a given year. In order to illustrate some of the potential implications, we write down an endogenous growth model with and without the customer margin. In the context of this model, we find that the customer margin dramatically increases the size and growth contribution of the largest firms, but lowers the aggregate growth rate by diverting resources from research to customer acquisition activities.

---

\*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>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 Zakrajšek (2017) argue that such frictions play a role in business cycle fluctuations,

Eaton, Eslava, Jinkins, Krizan and Tybout (2021) present evidence and a customer search model of exporting firm dynamics. Bernard, Dhyne, Magerman, Manova and Moxnes (2019) document the importance of customers for Belgian inter-firm transactions. Bornstein (2018) argues that consumer aging interacts with customer inertia to explain the decline in both labor's share and firm entry in recent decades. Bagwell (2007) surveys models and evidence on the role of advertising in reaching and attracting customers.

In this paper, we use Visa debit and credit card transactions from 2016–2019 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> If Visa's 60% share is representative of all debit and credit card spending, then Visa spending patterns are relevant for around 40% of all 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

---

<sup>1</sup>The sample is anonymized. We can link transactions of the same card through a masked card identifier, but 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>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>).

transaction. We find that the number of customers dominates the decomposition across merchants and 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. For only the largest merchants does the store margin play a big role. The importance of customers is remarkably consistent across all retail categories, such as furniture, electronics, restaurants, or gas stations.

Our decomposition does not distinguish between adding low-spending vs. high-spending customers. If expanding stores and merchants tend to add low-spending customers, this will tend to overstate the contribution of new customers and understate the role of spending increases by retained customers. To address this, we show that retained customers do tend to increase their spending more at fast-growing merchants and stores. Even with a generous adjustment for the spending of gained and lost customers versus retained customers, however, we find that the extensive (customer) margin accounts for 60% of sales growth variation across merchants and stores.

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 a given year. This is consistent with a stream of results on the role of fast-growing firms in aggregate job creation, such as Decker, Haltiwanger, Jarmin and Miranda (2016). We find that most of this tail behavior in the Visa data reflects adding or losing customers.

Though in retail rather than the manufacturing, our evidence of a large extensive margin for customers is in the spirit of findings by Foster, Haltiwanger and Syverson (2008, 2016) and Hottman, Redding and Weinstein (2016). These studies estimate 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 and variety of their products. Baker, Baugh and Sammon (2020) also analyze customers using debit and credit card transactions, specifically from 2010 to 2015. Their focus is on 550 firms, 420 of whom are publicly traded and hence have observable stock returns. Their analysis, like ours, emphasizes the importance of the customer margin. Afrouzi, Drenik and Kim (2020) and Argente, Fitzgerald, Moreira and Priolo (2021) explore the relationship between price-cost markups and margins of consumer demand.

In order to illustrate some of the potential implications, we write down an endogenous growth model with firm dynamics that explicitly incorporates the customer extensive margin. In the context of the model, firms invest in improving the quality of their products each period, which generates growth in the aggregate. Innovation outcomes are stochastic, so firms are heterogeneous in their quality levels and growth rates. Firms spend on marketing to access customers each period. Because they sell more to each customer they access, firms with higher quality products spend more to access more customers. Customer acquisition thereby amplifies size differences stemming from quality differences across firms. Customers are a static function of current year marketing efforts, so firms have no incentive to lower markups and build their customer base dynamically. This is consistent with empirical evidence on Irish exporting firms in Fitzgerald, Haller and Yedid-Levi (2019).

After calibrating the model parameters based on the empirical facts we document, we show that customers dramatically amplify the effects of quality differences and therefore the market share and growth contribution of large firms. Yet, we also find that the customer margin actually undermines growth. This is because it diverts resources from research toward marketing.

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. Section 4 lays out the growth model, its calibration, and its quantitative properties. Section 5 concludes.

## 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 January 2016 and December 2019. 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>4</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 PIN-debit 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 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 merchant-specific 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>4</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 2016–2019 Visa data contain an annual average of 428 million cards, 31.5 billion transactions, and \$1.07 trillion in sales for the retail sector plus restaurants.<sup>5</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 2019 as 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 2016, going back to 2007, 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-2016 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 2016–2019, but we also report results that use larger merchants only for this longer panel of 2007–2019 (see [Online Appendix B](#)).

---

<sup>5</sup>[Appendix Table A2](#) provides these statistics for each year separately.

### 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:

$$S = N \cdot \frac{V}{N} \cdot \frac{S}{V}, \quad (1)$$

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:

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

To operationalize this decomposition, we take logs of both sides in equation (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 different merchants. In this broad sample, the customer margin accounts for 71% of sales variation across merchants, transactions per card account for around 4%, and the ticket size accounts for the remaining 25%. When we look at only online

transactions (Panel B), the customer margin falls to 70% of variation in online sales across merchants. In contrast, the customer margin accounts for 75% of variation in offline sales across about 1.8 million merchants in 2019. Of this 75%, 66% comes from cards per store and only around 9% 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 12,000 large “named” merchants, which Visa tracks all the way back to 2007, the store margin is much more important, accounting for 41% of the variation in sales vs. only 46% that is accounted for by cards per store.

In Table 2 we focus on the offline retail (plus restaurants) sector, now showing additional types of variation. The first row (Panel A) reproduces the corresponding 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 four years of data (2016–2019), 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 almost 3.9 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 (68.1%) is attributed to the changing number of cards per store, and the rest (16.4%) 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



Table 1: Sales Decomposition for Different Merchant Samples

	Stores	Cards/Store	Trans/Card	Dollar/Trans
<b>A. All Data</b>		0.711	0.040	0.249
( <i>N</i> = 2, 291, 382)		(<0.001)	(<0.001)	(<0.001)
		[ 0.718]	[ 0.172]	[ 0.482]
<b>B. Online</b>		0.697	0.066	0.236
( <i>N</i> = 603, 373)		(<0.001)	(<0.001)	(<0.001)
		[ 0.689]	[ 0.175]	[ 0.373]
<b>C. Offline</b>	0.094	0.660	0.034	0.212
( <i>N</i> = 1, 878, 903)	(<0.001)	(<0.001)	(<0.001)	(<0.001)
	[ 0.113]	[ 0.721]	[ 0.215]	[ 0.488]
<b>D. Offline Retail</b>	0.098	0.689	0.036	0.177
( <i>N</i> = 939, 304)	(<0.001)	(<0.001)	(<0.001)	(<0.001)
	[ 0.123]	[ 0.761]	[ 0.231]	[ 0.547]
<b>E. Offline Retail, “named”</b>	0.405	0.456	0.036	0.103
( <i>N</i> = 12, 079)	( 0.003)	( 0.003)	( 0.001)	( 0.002)
	[ 0.623]	[ 0.672]	[ 0.284]	[ 0.500]

Note: *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 parenthesis, 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.

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 2016–2019 data, aggregating variables at the store-year level (we have 7.7 million such observations), and adding store and year fixed effects. The customer 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 2016 to 2019, the number of unique customers explains the vast majority (80% or more) 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-priori 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 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 2016 to 2019), 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 70% 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 a greater number of cards at merchants or stores with higher overall sales. For Visa credit cards from 2016

Table 2: Decomposing Sales in Offline Retail

	Stores	Cards/Store	Trans/Card	Dollar/Trans
A. Across Merchants ( $N = 939,304$ )	0.098 [ 0.123]	0.689 [ 0.761]	0.036 [ 0.231]	0.177 [ 0.547]
B. Within Merchants over Time ( $N = 3,846,501$ )	0.164 [ 0.804]	0.681 [ 0.975]	0.098 [ 0.942]	0.056 [ 0.971]
C. Across Stores within Merchants ( $N = 1,926,714$ )		0.841 [ 0.972]	0.077 [ 0.809]	0.082 [ 0.933]
D. Within Stores over Time ( $N = 7,746,218$ )		0.817 [ 0.995]	0.134 [ 0.969]	0.049 [ 0.987]

Note: 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 2016–2019 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.

Figure 1: Decomposing Merchant Sales Growth by Industry



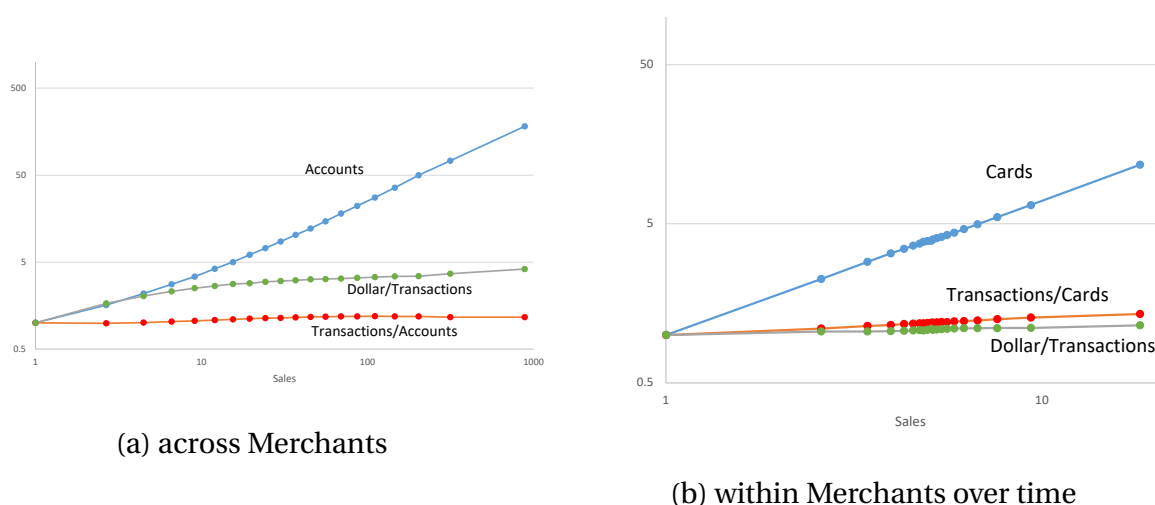
Note: This figure displays the coefficients of the “Within Merchant over time” decomposition by industry. The regressions are run with Visa data from 2016 through 2019, and include merchant and year fixed effects.

to 2019 we can match cards to households for about 50% of cards. These households 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 Table B2 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 customer 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

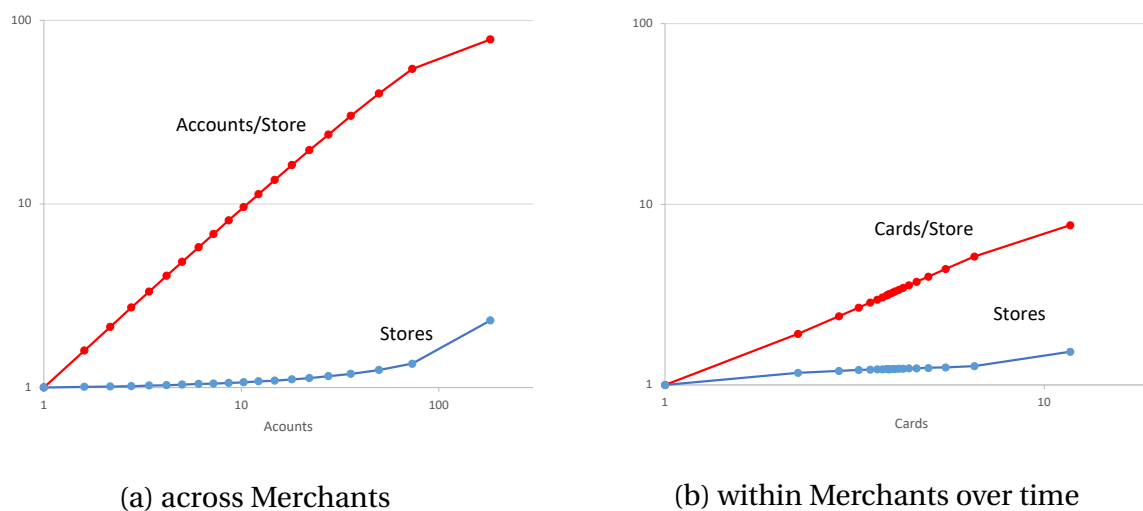


Note: Panel (a) is based on a cross section of all merchants in 2019. In panel (a), we group the x-axis into 20 bins, and report averages by bin, normalizing each variable by its average for the first bin. Panel (b) repeats the same exercise, but for the panel of merchant-years over 2016–2019. For (b) 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 unique customers into the number of stores and the number of unique cards per store, respectively. It shows that, both in the cross section and over time, the number of stores is not an important source of sales variation for the bottom half (in terms of sales) of merchants, which may be natural as many smaller merchants only have a single store. For larger merchants, stores become more important, in particular for the largest set of merchants (top ventile). This is similar to the role of estab-

ishments 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.

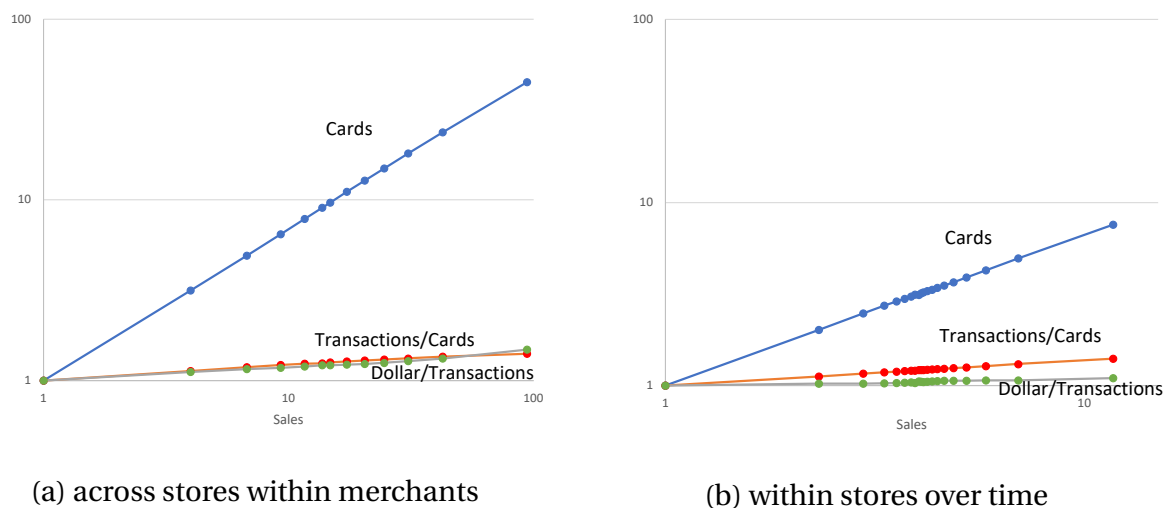
Figure 3: Stores vs. Cards Per Store



Note: 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 2016–2019. 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 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.

Figure 4: Decomposing Store Sales



Note: 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 2016–2019, de-meaning each variable by its store average and its year average. All panels are plotted on (base 10) log scale.

**Results by store age.** In [Online Appendix B](#) we repeat this analysis for the much smaller set of large merchants who are linked back to 2007 in the Visa data. The results look qualitatively similar, with the exception that the number of stores is much more important across large merchants and, to a lesser extent, over time within large merchants.

An advantage of a longer panel is that we can look at whether store dynamics differ by firm age. In [Online Appendix C](#), we decompose the sources of store growth separately for stores in their first two years since entry, years 3-5, and stores that have been open for 6+ years. Customers remain the primary driver of revenue growth for all three age groups (77%, 81%, and 85%, respectively), but new stores rely more than established stores on the average transaction amount to grow their sales (15% vs. 6%, respectively).

**Returning vs. newly acquired customers.** One possible concern about the

above analysis is that it confounds compositional effects. For example, we might be overstating the extensive margin if returning customers increase their spending a lot at growing stores, but average spending does not grow much because new customers spend less than returning customers.<sup>6</sup>

To address this concern, we regress the log change of spending per returning customer on the log change of total sales for merchant-years from 2016 to 2019 (adding year fixed effects).<sup>7</sup> We report the coefficient for all retail (first bar) and separately by 3-digit NAICS in Figure 5. By this metric, returning customers account for 38% of the variance of sales growth in all NAICS. Their contribution ranges from 26% among clothing stores to almost 47% among food and beverage stores. This 38% is notably higher than the approximately 20% variation in sales that we attributed to the intensive margin earlier, when we did not adjust for composition. Still, we continue to find that the extensive margin accounts for most of sales growth variation (62%) by this metric. In [Online Appendix D](#) we report similar results at the store level.

### 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.<sup>8</sup>

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

---


<sup>6</sup>In the case of Pareto distribution of spending across customers, one might see the entry of new customers exactly offset the growing spending of returning customers.

<sup>7</sup>Unlike the earlier decomposition analysis, this is not a precise decomposition because – due to turnover of cards – tracking returning vs. new customers requires us to limit attention to the subset of cards that are active over two consecutive years.

<sup>8</sup>Since the volume of transactions on the Visa network has been steadily increasing over time, throughout this section we measure both aggregate and merchant sales in 2012 CPI dollars and re-scale each of them by Visa's share of the debit and credit card market by dollar volume in the corresponding year (obtained from <https://wallethub.com/edu/cc/market-share-by-credit-card-network/25531>). As mentioned in Section 2, in this part of the analysis we also excluded gasoline sales.



**Figure 5:** Spending per returning customer on firm sales growth



The figure reports the coefficient in the regression of annual log change of spending per returning customer on annual log change of total sales. An observation is a merchant-year level. The regression uses 2016-2019 data and includes a year fixed effect.

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 6 plots the contribution of each group to aggregate sales increases or decreases, averaged across the three observations 2016-2017, 2017-2018, and 2018-2019. 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 about 90%, and the top and bottom 25% contribute more than 99%. The patterns appear to be fairly symmetric for growing and shrinking merchants.<sup>9</sup>

As we noted, in Figure 6 the grouping of firms is done year by year. This implies that the identity of tail firms is changing from year to year. How important are *cumulative* sales increases and decreases to the aggregate increases and decreases from 2016 to 2019? To find out we rank merchants based on their cumulative sales changes from 2016 to 2019. This includes entrants among the growers and exiting merchants among the shrinkers.

Figure 6: Contribution to Aggregate Sales Changes

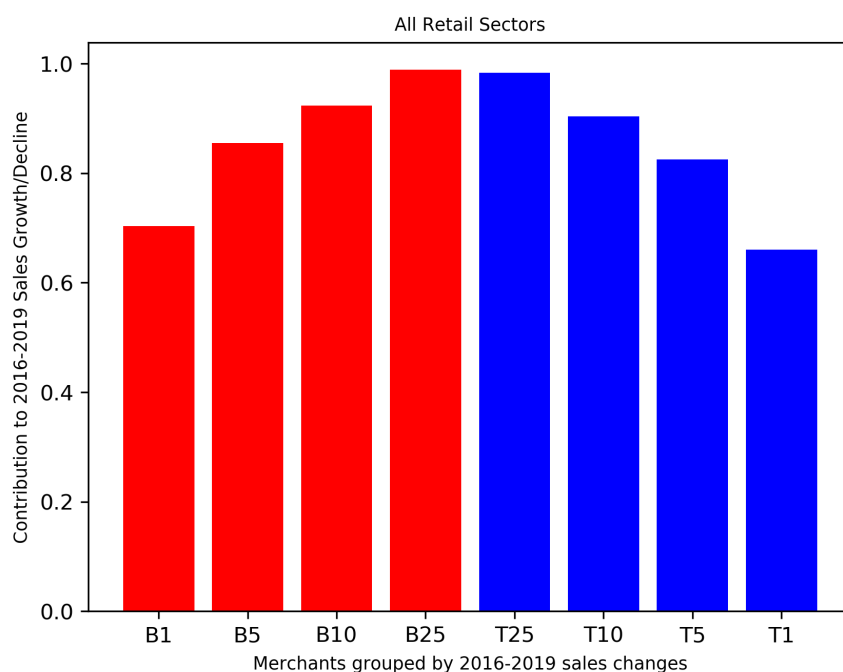


The figure reports the average contribution of each merchant group as defined in the text to aggregate sales change over year with the error bar extending one standard deviation up and down. An observation is a merchant-year and the figure uses a panel of merchants from 2016 to 2019. Each bar corresponds 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.

<sup>9</sup>In Figure B10 in the Appendix we show the share of initial sales at the top growers and shrinkers. It is notably smaller than their share of changes. For example, only 45% of sales on average are at the top 1% of growing firms.

In Figure 7, we then show that tail firm contributions remain remarkably similar when looking at cumulative changes from 2016 to 2019. Evidently, many firms are growing and shrinking by large amounts over the three 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 (2007–2019) panel, which include a much smaller set of merchants.


**Figure 7: Persistence of Merchant Contributions**



The figure reports the contribution of each firm group as defined in the text to aggregate sales change between 2016 and 2019. Each bar corresponds to a firm group. TX refers to top X% firms and BX refers to bottom X% firms by the absolute sales changes between 2016 and 2019.

Figure 8 reports the contributions of the top and bottom 1% of merchants for each 3-digit NAICS from 2016 to 2019. 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 NAICSs that extreme growers and shrinkers account for a large fraction of aggregate sales changes.

Figure 8: Contribution to Aggregate Sales Changes By NAICS



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 2016 to 2019. The calculation of each merchant group's contribution is described in the text.

**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 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

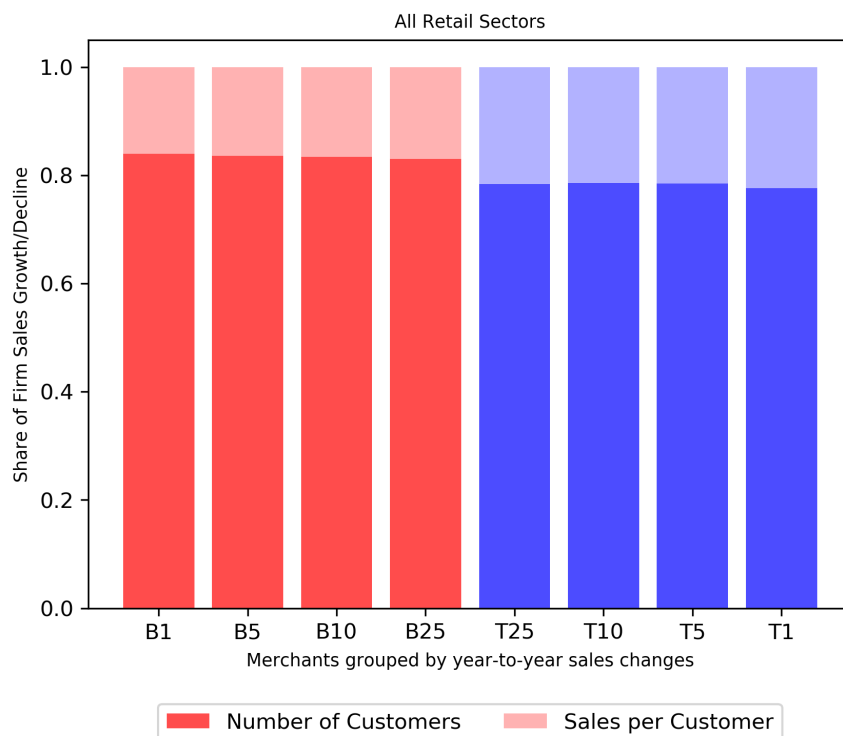
$$\overline{N}_{it} \equiv \frac{N_{it} + N_{i,t-1}}{2}$$

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 tell how much of the aggregate sales changes in each group are attributed to changes in the number of unique customers versus changes in sales per customers. Figure 9 shows that the change in customers accounts for around 80% of sales changes in the tails (modestly under 80% for increases, and modestly above 80% for decreases). 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 9:** Customers vs. sales/customer and firm sales changes



The figure reports the average share of sales changes in each firm group that can be attributed to changes in number of customers and changes in sales per customer respectively from 2016 to 2019. By construction, the two shares sum to 1. Each bar corresponds to a firm group. TX refers to top X% firms and BX refers to bottom X% firms by firms' absolute sales changes.

## 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. We should emphasize at the outset that the role of model in this paper is 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:

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

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

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

## 4.2. Firms

Each firm uses production labor  $l_{it}$  to produce its single variety:

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

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. As  $\gamma \rightarrow \infty$ , we obtain the polar economy where sellers reach all customers ( $n_{it} = 1$ ) with no need to expend marketing effort ( $m_{i,t} = 0$ ). We assume no negative externality with respect to other firms' marketing efforts.<sup>10</sup>

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

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

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

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

---

<sup>10</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 firms 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 E](#) contains the model with a parameterized marketing externality.

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

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

The firm'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 firm'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:

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

Defining a firm'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 firm 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  $x_{it} \in [0, 1]$ :

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

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 firm 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 firms, the less labor required to successfully innovate ( $\zeta > 0$ ). A firm's value function is given by:

$$V_t(z) = \pi_t(z) + \max_{x \in [0,1]} \{ R_t^{-1} [x V_{t+1}(ze^{\Delta-gt}) + (1-x) V_{t+1}(ze^{-gt})] - s_t(z, x) \}$$

where  $R$  is the gross interest rate. The Euler equation produces the usual relationship between the growth rate  $g$  and the consumer's discount factor in the absence of aggregate uncertainty:

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

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

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

#### 4.4. Labor market clearing

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

$$\begin{aligned} L_t &= \int l_t(z) \, dF_t(z) \\ M_t &= \int m_t(z) \, dF_t(z) \\ S_t &= \int s_t(z) \, dF_t(z). \end{aligned}$$

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:

$$L_t = \frac{\gamma(\theta - 1)(1 - S_t)}{\gamma(\theta - 1) + 1} \quad \text{and} \quad M_t = \frac{L_t}{\gamma(\theta - 1)}.$$

#### 4.5. Calibration

A period in the model is one year. Without loss of generality we set the aggregate labor endowment to 1.<sup>11</sup> 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

---

<sup>11</sup>The model features strong scale effects, but is invariant to the scale of the economy once re-calibrated to match the targeted moments such as the growth rate.

rate is set to 2.9% per year to match the data (see below), the steady state real interest rate is 6.7% per year as in Farhi and Gourio (2018).

- We set the level of marketing costs  $\phi$  so that the firm with maximum relative quality reaches half of the customers.<sup>12</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 firm.
- 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  $\Delta$  of 6%, roughly double the growth rate. As a consequence the average probability of innovation success ( $x$ ) is around one half. With  $\theta = 3$ ,  $\gamma = 1.25$ , and  $\Delta = 0.06$ , sales grow by 29% ( $g \cdot \xi_{y,q}$ ) for expanding firms and shrink by 29% for contracting firms.
- We choose the innovation cost parameters  $\lambda$  (the scale of R&D costs) and  $\zeta$  (the research spillover or convexity parameter) to achieve a 2.9% growth rate and for the top 1% fastest growing firms to account for 70% of sales changes. The growth rate is for labor productivity in the retail sector from 1988–2019 according to the U.S. Bureau of Labor Statistics. 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 firm.

---

<sup>12</sup>For confidentiality reasons we cannot target the precise share of customers reached by top firms in the Visa data.

Table 3: Parameter Values

Symbol	Parameter	Value
$\sigma$	Intertemporal elasticity of substitution	0.50
$\theta$	Elasticity of substitution between varieties	3.00
$\beta$	Discount factor	0.992
$\phi$	Scale of marketing costs	$3.57 \cdot 10^{37}$
$\gamma$	Elasticity of marketing costs wrt customers	1.25
$\Delta$	Quality step size	0.06
$\lambda$	Linear research cost	0.094
$\zeta$	Research spillover parameter	10.04

## 4.6. Results

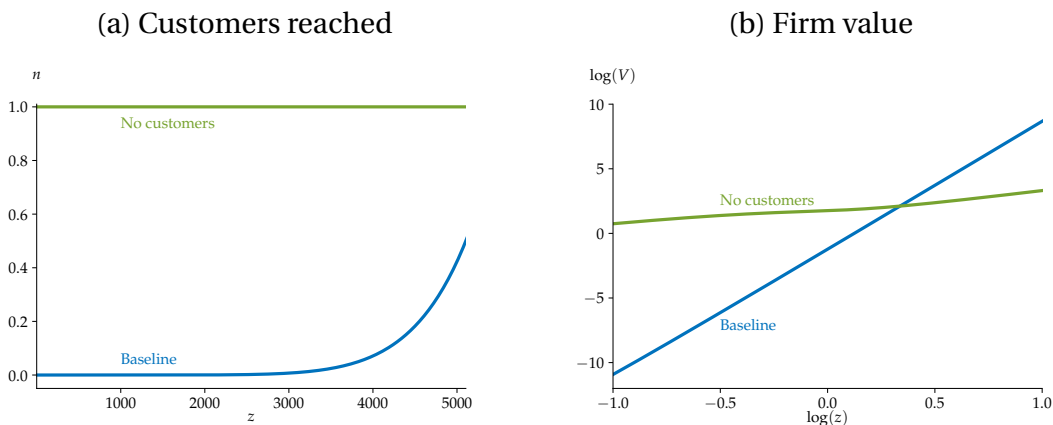
We are now ready to characterize equilibrium outcomes. For contrast we will show what happens in an economy with no customer margin. We achieve this by setting  $\gamma = \infty$  (and therefore  $\Gamma = 1$ ), so that labor is not needed to access customers (for all firms,  $n = 1$  despite  $m = 0$ ). We keep all other parameter values the same when we make this comparison.<sup>13</sup>

Figure 10a shows how  $n$ , the fraction of consumers the firm sells to, varies with the firm's relative quality  $z$ . It is log-linear with elasticity  $\gamma/(\gamma - 1)$  in the Baseline. This in turn makes the value of the firm much more convex with respect to  $z$  in the Baseline than in the No Customers case – see Figure 10b.

Because the customer margin makes variable profits increase much faster in relative quality, it induces higher quality firms to do more innovation than otherwise. This can be seen in Figure 11a. A corollary is that R&D intensity (research spending as a share of sales) is slowly decreasing with respect to  $z$  in the baseline case, whereas it falls quickly with  $z$  in the model without a customer

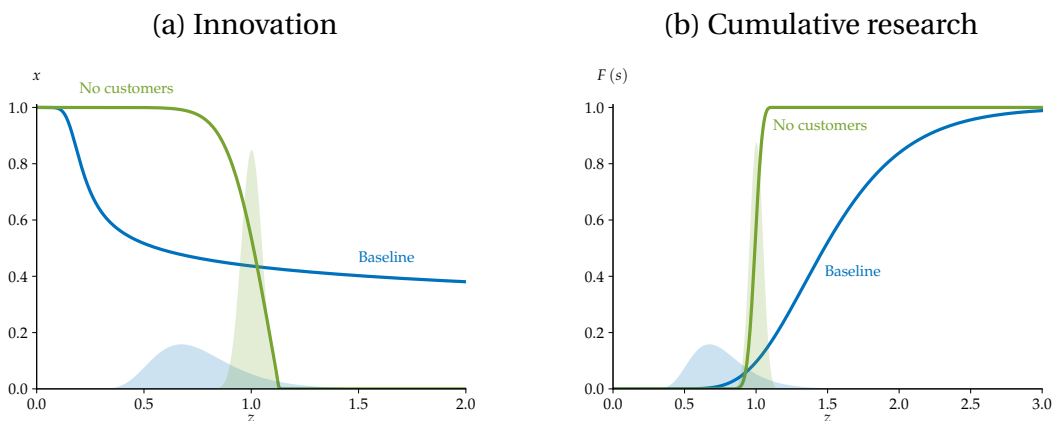
<sup>13</sup>With no customer margin, the elasticity of sales with respect to quality is only  $\xi_{y,q} = \xi_{c,q} = \theta - 1$ . Thus sales will be much less responsive to relative quality than in the economy with a customer margin.

Figure 10: Customers and firm value



Note: In panels (a) and (b), this figure shows how the fraction  $n$  of customers reached by the firm and its value  $v$ , vary with its relative quality  $z$ , respectively. The baseline features  $\gamma = 1.25$  and the “no customers” version uses  $\gamma = \infty$ .  $\gamma$  is the elasticity of marketing costs with respect to customers.

Figure 11: Innovation

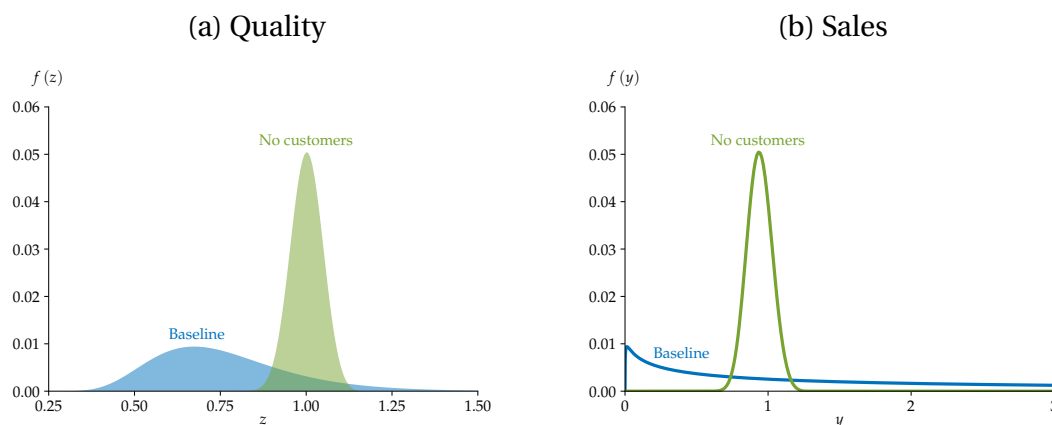


Note: This figure shows how the arrival rate of innovations  $x$  varies with the firm’s relative quality  $z$ . The Baseline features  $\gamma = 1.25$  and the “No customers” version uses  $\gamma = \infty$ , where  $\gamma$  is the elasticity of marketing costs with respect to customers.

margin. This shifts the bulk of R&D to large firms, as depicted in Figure 11b. In the baseline model, marketing and R&D expenditures are complements.<sup>14</sup>

Because higher quality firms do more R&D, the stationary distribution of relative qualities is much more dispersed when there is a customer extensive margin than when all firms reach all customers (Figure 12a).<sup>15</sup> In turn, the distribution of sales across firms is dramatically more dispersed when there is an extensive margin for customers (Figure 12b). Higher quality firms have more customers, and this endogenously induces more quality dispersion.

Figure 12: Distribution of quality and sales



Note: This figure shows the density of firm relative quality  $z$  in panel (a) and of firm sales  $y$  in panel (b). The Baseline features  $\gamma = 1.25$  and the “No customers” version uses  $\gamma = \infty$ , where  $\gamma$  is the elasticity of marketing costs with respect to customers.

In Table 4 we compare some variables in steady state across the Baseline and No Customer cases. The endogenous growth rate of aggregate quality actually rises from 2.9% in the baseline to 3.2% in the model without a customer margin. This flies against the intuition that the customer margin encourages firms to invest in quality innovations by making their profits more convex in quality. The countervailing force is a general equilibrium one: marketing labor steals labor

<sup>14</sup>See Cavenaile and Roldan-Blanco (2021) for a theoretical and empirical analysis of complementarity versus substitutability between R&D and marketing expenditures.

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

away from research (and production). Fully 27% of all labor is freed up from doing marketing when going from the Baseline to the No Customers model. Production labor soars from 68% of all labor in the Baseline to 93% with no customer margin. But there is still room for research labor to rise from 4.9% in the baseline to 6.7% of all labor in the model with no extensive margin for customers. The third column of Table 4 scales down the labor endowment in a No Customers economy so that there is no more labor available for production and labor than in the Baseline economy. Here we see that the growth rate does indeed fall to 2.65% compared to the baseline level of 2.90%.

Table 4: Steady-state endogenous variables

Symbol	Variable	Baseline	No customers	No customers (scaled)
$g$	Growth rate	2.90%	3.22%	2.65%
$r$	Interest rate	6.70%	7.36%	6.19%
$L$	Production labor	67.9%	93.3%	93.2%
$M$	Marketing labor	27.2%	0.00%	0.00%
$S$	Research labor	4.93%	6.70%	6.83%
Labor endowment		1	1	0.728

Note: The Baseline economy feature the parameter values in Table 3. The No Customers economies use  $\gamma = 0$  but otherwise Baseline parameter values.

Interestingly, in the scaled economy the *fraction* of labor devoted to research is actually higher at 6.8% than in the baseline economy's 4.9%. Because the scaled labor endowment is 0.728 instead of 1, this implies that 0.0497 units of labor are allocated to research in the scaled No Customers economy, versus 0.0493 in the baseline model. Why does the scaled economy grow more slowly than the Baseline economy despite devoting more effort to research? The answer appears to be the benefits of having high quality firms improve their quality further in the Baseline economy.

Table 5 shows the importance of the covariance between sales shares and quality growth to aggregate growth. To illustrate this we calculate a Törnqvist approximation to the true growth rate:

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

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

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

The Table indicates that this Törnqvist second-order approximation is quite good. If we instead use initial sales shares as weights (just using the 1st order term, a 1st order approximation), the approximated growth rate understates the true growth rate by about 50 basis points in the case with a customer margin. With no customers, in contrast, the second order term is much smaller. This breakdown illustrates the importance of customer acquisition in amplifying the growth contribution of the right tail of firm growers. Firm size dispersion is much higher with customer acquisition, so right tail growers contribute much more to growth in the Baseline.

Just like in the data, we can calculate the contribution of the top 1% of firms to aggregate sales increases. As depicted in Figure 13, our baseline model is calibrated to achieve a contribution of 70% from the top 1%. Without a customer margin, in contrast, the top 1% of firms 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

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



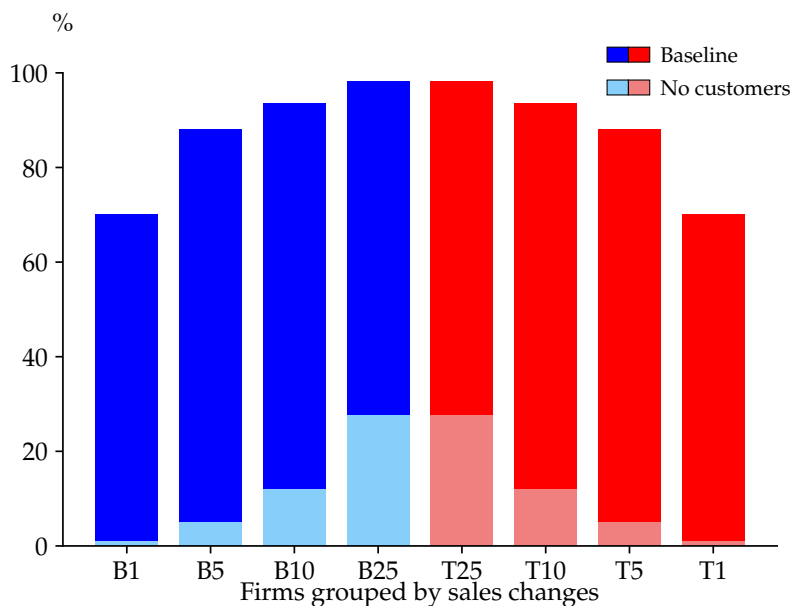
contribution of top firms is the same as their contribution to aggregate sales increases. Therefore, the top 1% fastest growing firms also account for 70% of consumption growth.

**Table 5:** Törnqvist growth decomposition

	Baseline	No customers	No customers (scaled)
True growth rate	2.90%	3.22%	2.65%
Approximated growth rate	2.95%	3.26%	2.70%
1st order term	2.49%	3.17%	2.61%
2nd order term	0.46%	0.09%	0.09%

Note: The Baseline economy feature the parameter values in Table 3. The No Customers economies use  $\gamma = 0$  but otherwise Baseline parameter values. The scaled economy has a total labor endowment of 0.728 instead of 1.

**Figure 13:** Firm Contributions to Aggregate Sales Changes



## 5. Conclusion

Using Visa data on credit and debit card transactions at U.S. retail merchants from 2016 to 2019, we document 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 write down a simple growth model that incorporates the customer extensive margin and illustrates how and why the customer margin may matter for growth. In the model, firms pay marketing and research costs to acquire customers and improve their quality. The customer margin makes large firms drastically more important for sales and overall growth. The ability to increase profits by adding customers increases the returns to innovation and stimulates research. But marketing diverts resources away from research. The latter force dominates so that growth is modestly lower in the model with a customer margin than if all firms could access all customers.

## References

- Afrouzi, Hassan, Andres Drenik, and Ryan Kim, “Growing by the masses: Revisiting the link between firm size and market power,” 2020.
- Argente, David, Doireann Fitzgerald, Sara Moreira, and Anthony Priolo, “How do entrants build market share? The role of demand frictions,” 2021.
- 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.
- Bagwell, Kyle, “The economic analysis of advertising,” *Handbook of Industrial Organization*, 2007, 3, 1701–1844.
- Baker, Scott R, Brian Baugh, and Marco C Sammon, “Measuring customer churn and interconnectedness,” 2020. 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,” 2019. NBER Working Paper 25441.
- Bornstein, Gideon, “Entry and profits in an aging economy: The role of consumer inertia,” 2018. working paper, Wharton.
- 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, pp. 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,” 2019. 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 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.

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), 2509–39.
- 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.