

ONLINE APPENDIX

Assessing the Gains from E-Commerce

by Dolfen, Einav, Klenow, Klopach, Levin, Levin and Best

February 2019

This online appendix contains the following three items. In Appendix A we discuss the main datasets that we use for our analysis. Appendix B provides further detail on how we measure online spending in our data and the definition of E-Commerce we use. Finally, Appendix C contains further details about derivations and sources for numbers that appear in the text.

Appendix A. Data and Samples

A.1. The Datasets

Visa Transaction Table

We combine several datasets for our analysis. The main dataset is a proprietary dataset by Visa Inc. covering the universe of transactions on the Visa network. This dataset is at the level of the transaction and contains transactions by both credit and debit cards. We observe transactions starting in 2007 and up to (and including) 2017.

The main variables for our analysis are the card number, a merchant identifier, the transaction ZIP code (available for brick and mortar transactions), the transaction amount and the transaction date. One limitation of this dataset is that we are unable to distinguish between different outlets of the same merchant in the same ZIP code. We will address this issue with a different Visa dataset (the GMR table discussed below). To that end, the transaction table also contains an establishment identifier (available since mid 2015) which can be linked to the GMR table.

We furthermore observe two variables that identify whether or not the card was present for the transaction (the card would be present for a brick and mortar transaction, but not for e.g. an E-Commerce transaction). The first indicator is always available and distinguishes between card present (hereafter referred to as CP) and card not present (hereafter referred to as CNP). The second indicator allows a further breakdown of CNP transactions into various categories, namely E-Commerce, Mail Order, Phone Order or Recurring Transactions (e.g. phone bills). These two indicators will be the basis for our measures of E-Commerce on the Visa network.

We do not directly observe any card attributes in this dataset. We however create a card-year location variable based on the brick and mortar transactions of the card-year. In particular, we define a card-year's location to be the transaction weighted average of the longitudes and latitudes of the brick and mortar transaction locations of that card. We only use transaction ZIP codes in which the card transacted 20 or more times in a given year to avoid contamination of the card location by e.g. transactions during holidays. We use the ZIP code centroid to assign longitude and latitude to transaction ZIP codes. We also have access to a different measure of card location based on credit bureau data (discussed below). This is however only available for a subset of cards and more recent years. It is worth noting that our measure of card location based on a card's transaction performs very well when compared to this external data.

One important limitation of the transaction data concerns the merchant identifier. While every transaction is assigned a merchant identifier, this identifier does not always allow us to infer the exact merchant. The Visa data distinguishes between two types of merchants, 'named' and 'unnamed' merchants. Roughly speaking, 'named' merchants are large chains for which Visa assigns a unique merchant ID, i.e. there is a one to one mapping between Visa's merchant id and the merchant. 'Unnamed' merchants are typically smaller chains and single establishment merchants. All 'unnamed' merchants within the same industry are assigned the same merchant id. As a consequence it is not possible to identify the actual merchant behind a Visa merchant id for 'unnamed' merchants. We will at

times restrict our analysis to named merchants for the parts of the paper for which identifying the exact merchant is important. 58% of dollars in our sample are transacted at named merchants.

There are additional merchant variables in the original dataset, such as a “merchant string”. This is the merchant name that would e.g. appear on a credit card statement. While this could in theory allow us to a.) distinguish between smaller (unnamed) merchants that carry the same merchant id and b.) disentangle different stores of the same merchant in the same ZIP code or, the merchant strings are in practice very fuzzy and cannot be easily linked.

Visa GMR Table

Global Merchant Repository (GMR) is an effort by Visa to create a master file of merchant information from data provided by the merchant’s acquiring bank and from external data providers. All Visa transactions were linked to a GMR entity via a unique identifier starting in mid 2015. Each GMR-stamped transaction is mapped to a merchant ID and a store ID. For each store, GMR contains the mailing address and the corresponding latitude/longitude pair.

Credit Bureau Data

We have access to an additional dataset that provides cardholder-level demographics which can be linked to a sample of Visa credit cards. This dataset is provided by a large credit bureau. About 50% of active credit cards in 2016 and 2017 were linked to an entity in this dataset. An additional 7% were linked to multiple rows in the dataset; we discard these records. For each cardholder matched to the credit bureau data, we observe the cardholder’s age and their 9-digit billing ZIP code, as well as their estimated household income, marital status, number of children, and education level.

A.2. The Sample

The Transaction Sample

We impose several sample restrictions. We focus on the transaction of Visa credit and debit cards (i.e. we discard non-Visa cards as well as Visa Pre-Paid Cards) at U.S. merchants. Furthermore, we focus on Credit and Debit-Signature transactions only. This mainly excludes Debit-PIN transactions. Following the Durbin Amendment in 2010 (part of the Dodd-Frank Bill), Visa was not able to restrict how merchants routed Debit-PIN transactions. Therefore, starting in 2012 (when the law went into effect), the data exhibits significant fluctuations in the Debit-PIN transactions of stores. One day or hour a store is transacting with Visa and the next day it looks like they have 0 transactions and the next day they are back again. All the while, their neighbors stay steady on Visa. We hence focus on Credit and Debit-Signature transactions where merchants’ network routing is fairly consistent. Transactions worth 91.5% of total dollars on the Visa network satisfy these filters.

We furthermore impose the additional restriction that cards in our sample must have transacted with at least five merchants over their lifetime. This filter was chosen to exclude cards that are only used for one merchant and gift cards (there is a large number of cards that only transact with one merchant for a total of USD 50 or USD 100). Transactions worth 87.6% of total dollars satisfy all filters combined.

The Convenience Sample

In the convenience analysis, we use 2017 transactions from a 10% random sample of the cards that were matched to the credit bureau data in five “mixed” retail NAICS categories – i.e., those that had online share between 10% and 90%. Those include the following 3-digit NAICS codes: 442, 443, 448, 451, and 453. We count offline transactions as those that were marked as occurring face-to-face (i.e. with the card physically swiped) and online as those that were marked with the E-Commerce indicator. We exclude phone order, mail order, and recurring transactions from this analysis.

For each transaction, we calculate the distance between the merchant and the card as the distance between the card’s ZIP+4 from the credit bureau data and the closest offline branch of that merchant (defined as the latitude and longitude of the store as recorded by Visa). We keep transactions that occurred at merchants that had an offline presence within 50 miles of the consumer’s location. We also exclude merchants that had a greater than 99% online share or less than 1% online share within our sample transactions.

The Variety Sample

The variety sample consists of a random 1% sample of cards in 2017. We only consider transactions of these cards at named merchants (because controlling for the exact merchant identity is important) and in the narrower set of E-Commerce industries (excl. Nonstore Retail) (3-digit NAICS 441, 442, 443, 444, 445, 446, 448, 451, 452, 453). We choose these industries (as opposed to the baseline set of industries) because our estimation strategy again relies on distance to the merchant which is less relevant in the Non-Retail E-Commerce industries (e.g. Hotels, Car Rental). We also exclude Nonstore Retail because distance to brick and mortar stores is not meaningful in Nonstore Retail. We furthermore restrict our analysis to transactions for which the card is located within 20 miles of the merchant (using the transaction based measure of card location and locating a physical store and the ZIP code centroid).

The first σ we estimate is based on choices between online and offline merchants. For that we use all the transactions in the Variety Sample. We construct all pairs of physical store j and CNP merchant k such that card i buys from one of these. We require that both merchants are in the same 3 digit industry, that the store is within 20 miles of the card location and that merchant k has CNP revenues in that year.

The second estimated σ is based on the comparison of different offline choices. For this estimation, we only use the CP transactions in our Variety Sample. We then construct, for each individual i and 3-digit NAICS, all pairs of physical stores j and k such that i buys in at least one of these stores. We furthermore require that both stores are in the same 3 digit industry and both within a 20 mile radius of the card location.

Appendix B. Measuring E-Commerce in the Visa Data

E-Commerce Variables in the Visa data

In the following we will distinguish between transactions that are CNP (Card Not Present) and CP (Card Present). CP transactions are brick and mortar transactions whereas CNP transactions refer to the aggregate of E-Commerce, Recurring Transactions (e.g. utilities, phone bills), Mail Order and Telephone Order.

As highlighted in the Data section, we observe two variables at the transaction level that allow us to distinguish whether a card was present or not during the transaction. The first is a CNP indicator that distinguishes between CP and CNP transactions. This variable is automatically created by Visa and is available for every transaction. The second variable is an E-Commerce indicator. This variable is filled in by merchants and allows a further breakdown of CNP transactions into E-Commerce, Mail Order, Telephone Order and Recurring Transactions. One limitation of this latter indicator is that it is not required by Visa and 40% of values are missing. It is worth noting that the share of missing values has been declining over time. Furthermore this indicator also contains ambiguous values for 14% of transactions. This implies that we often cannot identify whether a CNP transaction is E-Commerce or not. Since this variable is filled in by merchants, whether or not the E-Commerce indicator is missing varies from merchant to merchant.

In our final sample, 47% of transaction dollars are classified as CNP using the CNP indicator. Furthermore 20% of dollars can be classified as E-Commerce and 4% as either Mail Order, Telephone Order or Recurring Transaction using Visa's E-Commerce indicator. This implies that half the CNP transaction dollars on the Visa network cannot be broken down any further using information contained in Visa's datasets.

As a consequence we will not estimate E-Commerce in the Visa data by only using Visa's E-Commerce indicator. First, the large number of missing values would bias downwards our estimates of the E-Commerce share. Second, the declining share of missing values over time would bias upwards our estimate of the rise in E-Commerce. We will also not classify E-Commerce by only using the CNP indicator either because this measure would overestimate the E-Commerce share since it also includes other CNP transactions. In the following we will discuss how we combine the information of both these indicators to obtain our estimate of E-Commerce in the Visa data.

Measuring E-Commerce in the Visa Data

Given the data limitations, we combine the information of both the CNP and E-Commerce indicators to estimate E-Commerce spending. The underlying idea is to create (3-digit industry-year) weights that map CNP spending into E-Commerce spending. We choose this strategy as every transaction on the Visa network is either classified as CP or CNP.

To create these weights, we only keep CNP transactions that have a valid E-Commerce indicator (i.e. allow us to distinguish whether or not the CNP transaction was an E-Commerce transaction). We then calculate the share of CNP dollars that are E-Commerce dollars on this clean subsample. We do this exercise by 3-digit industry year to allow for different mappings across industries and time. To then obtain an estimate of E-Commerce spending by

industry-year, we multiply CNP spending in that industry-year in our full sample by the weights we calculated in the previous step. We find that 40% of transaction dollars are classified as online dollars using this methodology (Recall that 47% were CNP).

The E-Commerce Industries

The procedure described in the previous subsection yields estimates for online dollars on the Visa network in each industry-year. We will however not include the online dollars from all industries when calculating our final measure of E-Commerce on the Visa network. This decision is based on our definition of E-Commerce.

We define E-Commerce industries to be industries that are affected through gains in buying/ shopping convenience and/or increased variety by the rise of E-Commerce. We hence choose to not include industries in which the convenience is only in terms of payment. Examples of industries which we believe to only be affected through convenience in terms of payment are utilities, telecommunication and broadcasting.

Our baseline set of E-Commerce industries is as follows: Motor Vehicle and Parts Dealers (441), Furniture and Home Furnishing Stores (442), Electronics and Appliances Stores (443), Building Material and Garden Equipment and Supplies Dealers(444), Food and Beverage Stores (445), Health and Personal Care Stores (446), Clothing and Clothing Accessories Stores (448), Sporting Goods, Hobby, Musical Instrument, and Book Stores(451), General Merchandise Stores (452), Miscellaneous Store Retailers (453), Nonstore Retailers (454), Air Transportation (481), Transit and Ground Passenger Transportation (485), Rental and Leasing Services (532), Administrative and Support Services (561), Accommodation (721)

We will however also report an alternative estimate of E-Commerce which is based on a narrower set of industries, namely the subset of the above industries belonging to the Retail NAICS (44 and 45).

Our Estimate of E-Commerce Spending on the Visa Network

To obtain our estimate of E-Commerce spending on the Visa network in any given year, we add the estimated online dollars across our E-Commerce industries and then divide by total spending on the Visa network in that year. As discussed above, we report two separate estimates, one using our baseline definition of the E-Commerce industries and one counting only the online dollars from our E-Commerce industries that are part of the Retail NAICS (44 and 45).

We classify 20% of the dollars spent on the Visa network as E-Commerce, and 13% when using the narrower set of industries. Recall that 47% of dollars are CNP and 40% are estimated to be spent online (without using our E-Commerce industry restrictions).

Appendix C. Figures and Tables

- Figure 1: Visa spending by year is calculated as total sales draft transaction spending on Visa credit and debit cards in our sample. GDP and Consumption estimates are from the U.S. Bureau of Economic Analysis (BEA).
- Table 1: List of 3-digit NAICS that we associate with E-Commerce, along with example merchants falling in each of the NAICS.
- Table 2: Estimate of the share of online spending on the Visa network in select 3-digit NAICS categories.
- Figure 2: Estimates of E-Commerce spending in the U.S. as a share of all consumption. We estimate E-Commerce spending on the Visa network (as discussed in Appendix B) and extrapolate it to the the U.S. economy assuming: 1) that Visa is representative of all card spending in terms of online share, and 2) all online spending is done using credit or debit cards. We first calculate, respectively, total amount of credit and debit card spending on the Visa network (by year) based on our transaction data. We then use external information on Visa’s share of total credit and debit card spending to calculate the total amount of card spending (by year) in the U.S.¹ Having calculated total card spending by year, we then multiply this by the corresponding online share on the Visa network. Using our two assumptions, this number is our estimate of total online spending in the U.S. by year. We then divide by total consumption to obtain our estimate of

¹The external information is provided by WalletHub: <https://WalletHub.com/edu/market-share-by-credit-card-network/25531/>. WalletHub calculates market shares for credit and debit card spending based on the SEC filings of all major card providers.

the U.S. online share.² ‘All online’ refers to our baseline estimate of E-Commerce spending in all consumer categories. ‘Retail online only’ refers to our alternative estimate which only counts online spending in retail industries as E-Commerce. Total consumption (the denominator for each series) is from the BEA.

- Figure 3: This figure displays the 2017 online share in each county calculated from the Visa data and adjusted by the propensity of county residents to use a credit card. Each card is placed in a county-card income bin according to their home billing ZIP code and estimated household income. We compute the online share for each county-card income bin from their Visa credit card spending (as discussed in Appendix B) and then adjust for differential propensities to use credit cards. The adjustment is made in the following way: Using the Experian data, we count the different number of Visa cards in the county - card income bins. We then obtain the population equivalent, i.e. the total number of people in the different groups, using 2015 IRS data on the number of tax filers. The population is calculated as the number of single filers + number of married filers $\times 2$ + number of head of household filers + number of dependents. Combining these two numbers, we calculate an adjustment factor that maps the different subgroups in Visa to their population equivalent, namely

$$\alpha_{cy} = \frac{\text{Visa Cards}_{cy}}{\text{Population}_{cy}}$$

The adjusted online share is then calculated using

$$\widetilde{\text{Online Share}}_{cy} = \alpha_{cy} \cdot \text{Visa Online Share}_{cy}$$

In the final step we scale down the different $\widetilde{\text{Online Share}}_{cy}$ such that the aggregate of $\widetilde{\text{Online Share}}_{cy}$ matches our estimated total U.S. E-Commerce share. The plot shows the online share (aggregated across cardholders of different incomes) within each county.


- Table 3: The table shows summary statistics for the transactions used in the convenience analysis. The ticket size panel gives the average dollars per transaction for each NAICS and channel (online or offline). Distance to the nearest store is calculated as the as-the-crow-flies distance between a consumer’s location and the nearest offline branch of the merchant where the transaction was made. The first row in each of the bottom two panels contains the average ticket size or distance. The numbers below, in parentheses, are the 10th and 90th percentiles.
- Figure 4: The figure shows the share of transactions that occur online as a function of the distance between the card and the nearest outlet of the merchant. The sample includes transactions made by 1% of cards in 2017 at merchants in the five mixed-channel NAICS listed in the data section. We include transactions at merchants that had a location within 50 miles of the card’s billing ZIP code. The black line shows a bin scatter of the share of these transactions that occurred online in the raw data. Each point gives the average share of transactions that were online for cards in a bin of size one mile. For example, the leftmost point on the black line shows that cards that were between zero and one mile away from an outlet of a merchant conducted about 12% of their transactions with that merchant in the online channel. The grey line shows the predicted share of online transactions from a logit regression of an indicator for whether the transaction was online on the distance between the card and merchant and a set of merchant fixed effects.
- Table 4: Each cell in the table gives the share of total online spending in 2014 by the amount of offline and online dollars spent at a given merchant by a card. Each observation in the underlying data is a card-merchant combination with an entry for offline and online spending. For example, the cell in the first row and third column contains the share of online dollars corresponding to card-merchant combinations where a card spent \$0 offline at a merchant and between \$10 and \$100 online at that same merchant. The "total" row (column) gives the sum of the cells across all columns (rows) in that row (column). All cells (excluding the total row and column) sum to 1.
- Table 5: Each column represents a separate regression. The estimates of ϕ are from the OLS regression $\ln M = \alpha + \frac{1}{\phi} \cdot \ln(oM_o + bM_b) + \epsilon$, where M denotes distinct merchants visited and $oM_o + bM_b$. One observation is a card-year. We run this regression separately for 2007 and 2017. As a robustness check, we ran this regression controlling for household income using credit reporting agency data. The sample is 127 million cards in 2017. For given card spending, richer households purchased from fewer merchants (elasticity -0.05). But the implied ϕ fell very little, from 1.69 to 1.68, once controlling for income.

²We also did an analogous aggregation taking into account the different online shares for credit and debit cards. The results are qualitatively and quantitatively very similar.

- Figure 5: The graph is based on a 1% random sample of cards in 2017. The underlying observations are card-store-merchant triples such that the card transacted either offline at the store or online at the merchant (or both), the store is within 20 miles of the card, and the store and the merchant are in the same 3-digit retail E-Commerce industry. The x-axis is distance of the store from the card (in 1 mile bins). The y-axis is percentage of online transactions out of total transactions. We aggregate to the distance level by summing the online and offline transactions across card-store-merchant triples. Finally the share of transactions online is calculated as a function of the distance to the store and the observations are connected with a smoothed curve.

We also conduct a related analysis of card choices between two (offline) stores as a function of distance to the stores. In particular, for card-store-store triples, we calculate the share of transactions at the farther store as a function of the differential distance between the stores. This relationship is depicted in Figure A1. The graph is based on a 1% random sample of cards in 2017. The underlying observations are card-store-store triples such that the card visited at least one of the two stores, both stores are within 20 miles of the card, and the stores are from merchants in the same 3-digit retail E-Commerce industry. The x-axis is differential distance of the two stores from the card (in 1 mile bins). The y-axis is the share of transactions at the farther store. We aggregate to the differential distance level by summing the farther and closer transactions across card-store-store triples. Finally the share of transactions at the farther store is calculated as a function of the differential distance and the observations are connected with a smoothed curve.

Figure A1: Relative trips as a function of distance



- Table 6: Each column represents a separate regression. Coefficients are from the regression $\ln\left(\frac{Trips_j}{Trips_k}\right) = \ln\left(\frac{q_j}{q_k}\right) - \sigma \ln\left(\frac{p_{jk} + \tau_j}{p_{jk} + \tau_k}\right)$. Observations are transactions from a 1% random sample of cards in 2017 wherein the card transacted with at least one of stores j and k at competing merchants in the same industry and in a retail E-Commerce NAICS category. In ‘online-offline’ j is a merchant with online sales and k a store within 20 miles of the card. In ‘offline-offline’ both j and k are stores within 20 miles of the card. The resulting tables at the card-merchant1-merchant2 level are then aggregated to a merchant1-merchant2-distance1-distance2 level where distance denotes store distance from the card (aggregated to 1 mile bins) by summing transactions. p_{jk} denotes the average ticket size across merchants j and k and τ a monetized cost of the return trip to the store. Both regressions are implemented using cross-store fixed effects (i.e., fixed effects for the (j, k) pair).

We also conduct several robustness checks for our estimate of substitutability. We focus on the ‘online-offline’ estimate because this is our baseline estimate for the welfare calculations. These baseline estimates (both in aggregate and by industry) are displayed in the first column of Table A1. We then estimate the same regressions using a 2% sample of the cards in 2017 for which we observe credit bureau data. The resulting estimates are displayed in the second column of Table A1. Lastly, we run the same regressions on this sample

Table A1: Substitutability Robustness

	Baseline	Credit Bureau Sample	Longitude-Latitude
Aggregate	4.3	5.8	6.3
Building Material, Garden Supplies	7.7	7.8	8.8
Motor Vehicle and Parts Dealers	7.5	7.3	8.2
Furniture and Home Furnishings Stores	7.4	7.6	7.5
General Merchandise Stores	5.8	6.2	6.3
Health and Personal Care Stores	5.5	5.3	5.9
Clothing and Clothing Accessories Stores	5.2	6.1	5.8
Miscellaneous Store Retailers	5.2	5.7	5.9
Sporting Goods, Hobby, Music, Book Stores	4.2	4.5	4.4
Food and Beverage Stores	3.6	5.9	6.6
Electronics and Appliance Stores	3.4	3.9	3.8

using alternative location measures. In particular, we locate cards using the longitude and latitude of their billing address (available from the credit bureau data) and locate stores using their longitude and latitude (available from the Visa GMR Table). The results of this regression are displayed in the third column of Table A1.

- Table 7: The consumption-equivalent welfare gain is $\left(\frac{1-s_{old}}{1-s_{new}}\right)^{\frac{\phi-1}{\phi} \frac{1}{\sigma-1}}$, where s denotes the U.S. online share in that year (holding Z , A_b and q_b constant). The results are obtained by substituting in the datapoints for s and using the values of ϕ and σ shown in the Table.
- Table 8: The income split is for the subset of households with credit reporting agency data on income. Counties are sorted by population density in 2017, then placed into top or bottom half of the population by density. County population is obtained from the 2010 Census.
- Table 9: Estimates are across offline versus online merchants within each listed NAICS category. For other E-Commerce NAICS categories (Air Transportation, Ground Transportation, Rental and Leasing Services, Administrative and Support Services, Accommodation) the offline component was sufficiently limited that we used the overall offline-online estimate of $\sigma = 4.3$. [By comparison, we tend to estimate higher elasticities of substitution between competing offline merchants. We estimate a pooled elasticity of 5.01 across offline retail merchants within 3-digit categories. Interestingly, for some categories with little or no online option, we estimate lower elasticities (such as 2.92 across restaurants, which is itself a 3-digit non-retail NAICS).]
- Table 10: We compare the welfare gains under nested CES preferences to our single nest benchmark. Each nest is a 3-digit NAICS. We distribute purchases at nonstore retailers (NAICS 454) to the other nests using eMarketer estimates of the composition of nonstore retail spending. The consumption equivalent welfare gain with nested CES preferences equals $\left(\prod_m (1-s_m)^{-\frac{\alpha_m}{\sigma_m-1}}\right)^{\frac{\phi-1}{\phi}}$. The results are obtained by substituting in the sector specific online shares s_m and elasticities of substitution σ_m . The outer nest Cobb-Douglas elasticities α_m are calibrated using spending shares. Note that we use sectoral online shares on the Visa network for this exercise. To account for fact that online spending is larger on the Visa network than in the overall economy we scale the resulting number down by multiplying it with the ratio of our baseline welfare estimates (using U.S. online shares) and the welfare estimate that results from using the online share on the Visa network instead. This can be thought of as a log-linear approximation.
- Table 11: Changes in online share are a sufficient statistic for assessing changes in spending per offline merchant, number of offline merchants visited and number of offline merchants in the market in our model (conditional on ϕ). The corresponding formulae are given by $b_{2017}/b_{2007} = [(1-s_{2017})/(1-s_{2007})]^{\frac{\phi-1}{\phi}}$, $M_{b,2017}/M_{b,2007} = [(1-s_{2017})/(1-s_{2007})]^{\frac{1}{\phi}}$, $M_{b,market,2017}/M_{b,market,2007} = (1-s_{2017})/(1-s_{2007})$. The results are obtained by using our baseline estimate of $\phi = 1.74$.

Appendix D. Estimates of consumer surplus including variety gains

1. The Consumer Problem: The first order conditions of the consumer problem are

$$\begin{aligned} o &= (\sigma - 1) \phi M_o^{\phi-1} F_o \\ b &= (\sigma - 1) \phi M_b^{\phi-1} F_b \\ \frac{o}{b} &= \left(q^{\frac{\eta}{\eta-1} \frac{\sigma-1}{\sigma}} \right)^\sigma \\ \frac{M_o}{M_b} &= \left[\left(q^{\frac{\eta}{\eta-1} \frac{\sigma-1}{\sigma}} \right)^\sigma \frac{F_b}{F_o} \right]^{\frac{1}{\phi-1}} \end{aligned}$$

The first order conditions pin down the online share s of the optimal consumption bundle, namely

$$s = \frac{oM_o}{oM_o + bM_b} = \frac{k}{k+1}$$

where $k = q^{\frac{\phi}{\phi-1}(\sigma-1)} \left(\frac{F_b}{F_o} \right)^{\frac{1}{\phi-1}}$. Furthermore it can be shown that

$$oM_o + bM_b = \frac{(\sigma - 1) \phi}{1 + (\sigma - 1) \phi} \times w$$

Using this, the relation between oM_o and bM_b and the identities $F_b = \frac{w}{A_b}$, $F_o = \frac{w}{A_o}$ we obtain the analytic solution to the consumer problem given in the main text.

2. Supply side: The optimal price of any firm can be shown to equal

$$p_m = \frac{\sigma}{\sigma - 1} \frac{w}{A}$$

This, combined with the free entry condition, pins down L_b and L_o to equal, respectively, $(\sigma - 1) K_b$ and $(\sigma - 1) K_o$. We then use the definition of shipping labor and the solution to the consumer problem to find

$$\begin{aligned} L_o &= \left(\frac{k}{k+1} \right) \left(\frac{1}{1 + (\sigma - 1) \phi} \right) L \\ L_b &= \left(\frac{1}{k+1} \right) \left(\frac{1}{1 + (\sigma - 1) \phi} \right) L \end{aligned}$$

Substituting the expression for production labor and shipping labor into the labor market clearing condition

$$M_{o,mkt} K_o + M_{b,mkt} K_b = \frac{1}{\sigma} \frac{(\sigma - 1) \phi}{1 + (\sigma - 1) \phi} L$$

Lastly, combining the zero profit conditions for online and offline merchants yields

$$\frac{bM_b/M_{b,mkt}}{oM_o/M_{o,mkt}} = \frac{K_b}{K_o}$$

Using the solution to the consumer problem then yields

$$\frac{M_{o,mkt}}{M_{b,mkt}} = k \frac{K_b}{K_o}$$

Combing this with the above expression of the labor market clearing conditions yields the analytic solution

$$\begin{aligned} M_{b,mkt} &= \frac{1}{1+k} \frac{1}{\sigma} \frac{(\sigma-1)\phi}{1+(\sigma-1)\phi} \frac{L}{K_b} \\ M_{o,mkt} &= \frac{k}{1+k} \frac{1}{\sigma} \frac{(\sigma-1)\phi}{1+(\sigma-1)\phi} \frac{L}{K_o} \end{aligned}$$

3. Estimating σ : The estimates of σ are based on the variety sample described in Appendix A.2. We restrict attention to transactions that are either CNP or CP and within 20 miles of (transaction based) card location. Based on this we will create two different datasets, each of which will yield a separate estimate of σ . The first dataset, hereafter referred to as ‘offline-offline’ dataset, is at the level of observation of card-store-store such that the card visits at least one of the stores. Both stores are required to be within a 20 mile radius of the card location and in the same 3 digit industry. The second dataset, hereafter referred to as ‘online-offline’ is at the level of observation of card-store-online merchant such that the card transacts with at least one of the entities. The store is again required to be within a 20 mile radius of the card location, the online merchant is a merchant with positive CNP sales that year and both are in the same 3 digit industry. In both datasets there are four additional variables, namely distance between the card and the merchant (set to zero for CNP purchases) (in 1 mile bins) and the number of transactions at each of the merchants. We then aggregate both datasets to the level of merchant j , merchant k , distance to j , distance to k by summing transactions and regress

$$\ln\left(\frac{Trips_j}{Trips_k}\right) = \ln\left(\frac{q_j}{q_k}\right) - \sigma \ln\left(\frac{p_{jk} + \tau_j}{p_{jk} + \tau_k}\right)$$

where p_{jk} is the average ticket size at merchants j and k (dollar weighted) and τ_j is the cost of travelling (a return trip) to j . τ_j consists of several components: First, we convert straight-line miles into driving miles (and driving time): 1 straight line mile requires 1.5 miles of driving on average (Einav et al, 2016), and one mile of driving requires 1.4 minutes of driving (Einav et al, 2016). Second, we calculate the time cost of driving. An average hourly after-tax wage of \$23 (BLS) implies a time cost of $1 \times 1.5 \times \frac{14}{600} \times 23 = \0.80 . Third, we calculate the monetary cost of driving. An average fuel plus depreciation per mile of \$0.53 (IRS) implies a monetary cost of $1 \times 1.5 \times 0.53 = \0.79 . Combining these three terms, the (round trip) cost of driving a (straight-line) mile is $2 \times (0.80 + 0.79) = \3.18 . We implement the above described regression using merchant cross fixed effects to control for $\ln(q_j/q_k)$. We run this regression on both datasets to obtain, respectively, the ‘offline-offline’ σ and ‘offline-online’ σ .

4. Consumer surplus: Denote by s the estimated U.S. E-Commerce share. It can be shown that welfare can be expressed as

$$W = \frac{(\sigma - 1)\phi}{[1 + (\sigma - 1)\phi]^{\frac{1}{\sigma-1}(\sigma - \frac{\phi-1}{\phi})}} A_b^{\frac{1}{\phi} \frac{1}{\sigma-1}} \times \left(\frac{1}{1-s}\right)^{\frac{\phi-1}{\phi} \frac{1}{\sigma-1}} \times \frac{w}{p}$$

Conditional on σ , ϕ and A_b , the consumption equivalent welfare gain Δ stemming from the rise in E-Commerce can be obtained from

$$\begin{aligned} W\left(\frac{w}{p}, s_{new}\right) &= W\left(\Delta \times \frac{w}{p}, s_{old}\right) \\ \left(\frac{1}{1-s_{new}}\right)^{\frac{\phi-1}{\phi} \frac{1}{\sigma-1}} \times \frac{w}{p} &= \left(\frac{1}{1-s_{old}}\right)^{\frac{\phi-1}{\phi} \frac{1}{\sigma-1}} \times \Delta \times \frac{w}{p} \\ \left(\frac{1}{1-s_{new}}\right)^{\frac{\phi-1}{\phi} \frac{1}{\sigma-1}} &= \left(\frac{1}{1-s_{old}}\right)^{\frac{\phi-1}{\phi} \frac{1}{\sigma-1}} \times \Delta \\ \Delta &= \left(\frac{1-s_{old}}{1-s_{new}}\right)^{\frac{\phi-1}{\phi} \frac{1}{\sigma-1}} \end{aligned}$$

Substituting in the discussed values for s , ϕ , σ will hence deliver the result. The welfare calculations by card income/ county density group are done analogously.³

5. Consumer surplus in nested CES case: The welfare gains in the nested CES case can be expressed as

$$\Delta = \left(\prod_m \left(\frac{1-s_{m,old}}{1-s_{m,new}}\right)^{\frac{\alpha_m}{\sigma_m-1}}\right)^{\frac{\phi-1}{\phi}}$$

where m denotes the nests, s_m the online share within nest, α_m the outer nest elasticity (Cobb-Douglas) and σ_m the nest specific elasticity. The α_m are calibrated using spending shares, the σ_m estimated by industry (analogously to the baseline σ) and the s_m observed in the Visa data.

³The underlying Visa E-Commerce shares for the different card groups are adjusted for the card-less as described above before substituting into the above formula.

6. Producer surplus/ Retail Apocalypse: Here we examine the impact of changing q_o/q_b , A_o/A_b on b , M_b , $M_{b, mkt}$ through the lens of the model. It can be shown that the online share s is a sufficient statistic for all three counterfactuals and that the predicted changes can be expressed as follows:

$$\frac{b_{2017}}{b_{2007}} = \left[\frac{1 - s_{2017}}{1 - s_{2007}} \right]^{\frac{\phi-1}{\phi}}$$

$$\frac{M_{b,2017}}{M_{b,2007}} = \left[\frac{1 - s_{2017}}{1 - s_{2007}} \right]^{\frac{1}{\phi}}$$

$$\frac{M_{b,market,2017}}{M_{b,market,2007}} = \frac{1 - s_{2017}}{1 - s_{2007}}$$

As we are describing the Retail Apocalypse we will use an estimate of the online share in the retail industries only, rather than in all of the U.S. economy. The online share in U.S. Retail is calculated analogously to the overall U.S. online share. In particular, we calculate, by year, total online revenues for online merchants in the retail NAICS. We then divide retail revenues on the Visa network by Visa's share of total card spending to obtain an estimate of total online spending at retail merchants in that year. In the final step we will divide this estimate by the BEA's Retail Trade Gross Output estimate to obtain an estimate of the online share in U.S. retail. The resulting estimates are 6.0% in 2007 and 9.5% in 2017.