THE AGGREGATE IMPACT OF ONLINE RETAIL*

by

Allen Tran†
UCLA

CES 14-23 May, 2014

The research program of the Center for Economic Studies (CES) produces a wide range of economic analyses to improve the statistical programs of the U.S. Census Bureau. Many of these analyses take the form of CES research papers. The papers have not undergone the review accorded Census Bureau publications and no endorsement should be inferred. Any opinions and conclusions expressed herein are those of the author(s) and do not necessarily represent the views of the U.S. Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed. Republication in whole or part must be cleared with the authors.

To obtain information about the series, see www.census.gov/ces or contact Fariha Kamal, Editor, Discussion Papers, U.S. Census Bureau, Center for Economic Studies 2K132B, 4600 Silver-Hill Road, Washington, DC 20233, CES.Papers.List@census.gov

Abstract

To study the impact of online retail on aggregate welfare, I use a spatial model to calculate a new measure of store level retail productivity and each store's equilibrium response to increased competitive pressure from online retailers. The model is estimated on confidential store-level data spanning the universe of US retail stores, detailed local-level demographic data and shortest-route data between locations. From counterfactual exercises mimicking improvements in shipping and increased internet access, I estimate that improvements in online retail increased aggregate welfare from retail activities by 13.4 per cent. Roughly two-thirds of the increase can be attributed to welfare improvements holding fixed market shares, with the remainder due to reallocation. Surprisingly, 8.2 percent of firms actually benefit as they absorb market share from closed stores. Finally, I estimate that the proposed Marketplace Fairness Act would claw back roughly one-third of sales that would otherwise have gone to online retailers between 2007-12.

[*] I wish to thank Hugo Hopenhayn, Pablo Fajgelbaum and Andy Atkeson for their guidance and encouragement, Shawn Klimek, Connan Snider and seminar participants at the Federal Reserve Bank of Richmond, Federal Reserve Bank of Philadelphia, Monash University, Queen's University and UCLA for useful comments, and Abigail Cooke for assistance at the California Census RDC. Any opinions and conclusions expressed herein are those of the author(s) and do not necessarily represent the views of the U.S. Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed

[†] Email: allentran@ucla.edu

1 Introduction

Probably the largest technological innovation of the past two decades is the internet, yet we know surprisingly little about its effect on welfare. A clear example of this is the development of online retail which has transformed the retail sector. Nominal sales at online retailers grew by an average of 17.5 per cent per year over the period 2000-13, compared to 3.3 per cent for the entire retail sector. As a consequence, online retailers' market share expanded from 2.7 to 17.8 per cent of total retail sales. The rapid adoption of online retail is obvious in the data yet the impact on economic welfare is unclear. In this paper, I study the impact of online retail, the embodiment of new internet technologies combined with improvements in logistics and shipping technology, on both retailers and consumers.

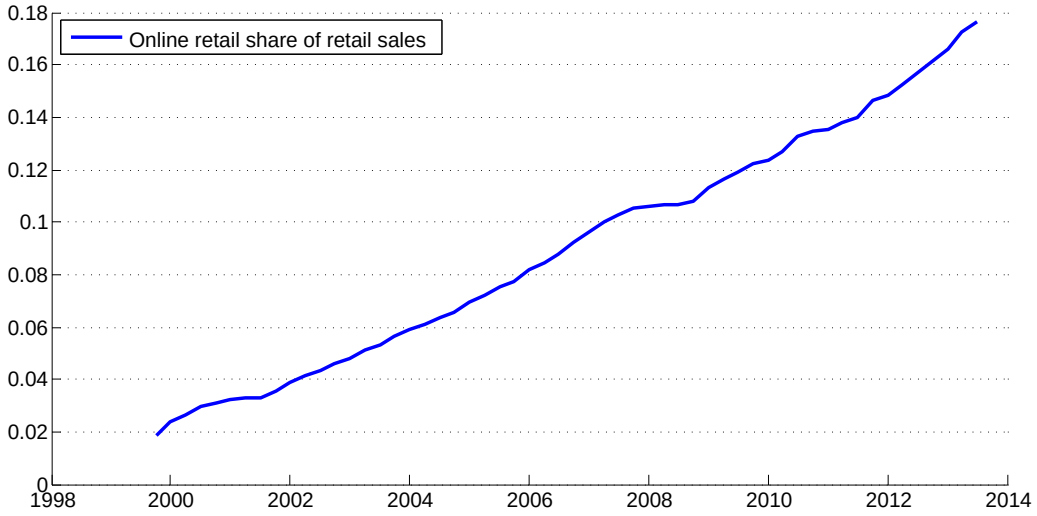

Figure 1: Growth of online retail sales

Understanding the impact of online retail is important for a number of reasons. First, the retail and wholesale trade sectors are large. Value added in the retail and wholesale trade sectors accounted for half of personal consumption expenditures on goods post 2000. Over and above interest in the retail sector, it offers insight into the effect of internet based innovations on the services sector, an increasingly common form of creative destruction driven by innovation from technology firms. Proposed legislation such as the Marketplace Fairness Act, a nationwide online sales tax, make understanding these issues particularly pertinent for policy.

There are two primary margins that matter when assessing the impact of online retail. First, consumers are likely to gain from both lower prices and the ability to shop at home, without incurring the cost of traveling to stores. Second, entry or improvements by online retailers alters the market structure of retail. Markups and sales per store fall, inducing store closures and a reconfiguring of the retailers participating in any local retail market. This extensive margin adjustment affects aggregate welfare through reallocation and also has ambiguous effects on consumer surplus as customers who are forced to shop at second-best options lose out.

Quantifying the impact of online retail requires the measurement these margins. However, both of these margins vary across markets since demand for online purchases varies across markets. Consumers without easy access to the internet and relatively high quality brick and mortar stores are less likely to seek online avenues to purchase goods. Therefore, an aggregate approach such as one based on a stylized two-sector model will not suffice. Here I explicitly account for geography and the heterogeneity of the US retail industry in order to accurately capture these margins.

In order to account for these margins properly, I estimate a spatial equilibrium model of retail with a geography based on the US retail industry. Stores exist in their real world locations and compete for customers described by zip code level demographic data. The model is estimated on confidential store-level data from the US Census of Retail Trade, demographic data from the Current Population Survey and Economic Census, and shortest-route data collated from Google Maps.

A central component of the model involves the measurement of the consumer surplus and profits, or total surplus, generated per purchase for each store. Total surplus per store reflects the potential welfare from a particular store. The actual surplus or welfare generated by a store depends on the allocation of consumers across stores.[1] When aggregated across stores, this measure captures the value of services generated in the process of transferring goods from retailers to final consumers, holding fixed the quantity of goods. To estimate store-level consumer surplus, I rely on the spatial nature of the model to separate effects due to the competition from nearby stores and the composition of customers, from a store's sales. This methodology can be applied to infer establishment-level productivities from other service industries where geographic proximity defines sets of competitors.

To quantify aggregate gains in welfare from online retail, I use counterfactual exercises which measure the effects from shocks to the equilibrium state. Counterfactuals isolate the causal effect from online retail as they hold other factors fixed by design. The counterfactuals are based on likely drivers of the increase in the online sales share from 2007-12. They include increases in internet access, reductions in shipping time and diversification into new industries by online retailers. Increased internet access and improvements in shipping that map to the data can account for roughly half for the observed increase in online retailers' share of sales. Although improvements by online retailers lead to a large amount of store closures, the gains in consumer surplus and the savings in store operating costs are enough to offset these losses. The 5.2 percentage point increase in online retailers' market share is associated with a 13.4% increase in aggregate welfare. While many traditional retail firms experience a reduction in profits, 8.2 per cent of traditional retail firms (inclusive of exiters) gain by absorbing part of exiting stores market share.

The final counterfactual quantifies the likely effects from the implementation of the Marketplace Fairness Act by imposing measures of state sales taxes on online purchases. I calculate the effect of the Marketplace Fairness Act by introducing the taxes after first projecting the model out from 2007 to 2012. The tax has a substantial effect in reallocating sales back to brick and mortar sales and mitigating store closures. Without the tax, sales at online stores more than double, leading to roughly 78,000 store closures over 2007-12. The tax has the effect of diverting one third of the lost sales back to brick and mortar stores preventing

[1] Prices divide the surplus between store and consumers and do not affect the size of the surplus from a purchase. However, prices do affect aggregate welfare through the allocation of consumers across stores.

18,000 store closures.

The paper is organized as follows. Section 2 describes the data that underlies most of the analysis in the paper. Details of the structural model are presented in Section 3. The description of the estimation routine for both the demand system and fixed costs and parameter estimates follow in Section 4. Results from the counterfactual exercises are reported in Section 5. Finally, I conclude in Section 6.

1.1 Related literature

1.1.1 Retail productivity

Both Foster et al. (2006) and Lagakos (2013) study retail productivity between "Mom and Pop" type stores and larger superstores using standard measures of productivity such as labor productivity and value added per worker. While these measures are informative at the aggregate level, they are less useful in evaluating store level retail productivity as store level versions of these measures are clouded by effects from competition. Markups are likely to be lower in relatively competitive markets, lowering value added per worker and hence estimates of productivity in the area. Because much of the reallocation caused by online retail occurs between stores, as opposed to firms, I use a different measure of productivity which neutralizes competitive effects. This measure of store-level retail productivity captures the consumer surplus net of costs from purchasing at a particular store, a measure that is independent of prices.

1.1.2 Reallocation and productivity growth

In the model, technological change comes in the form of entry or improvements by online retailers. Therefore, the driving force of technological change is reallocation rather than across the board improvements in productivity. This mirrors the literature on reallocation from technological improvements for manufacturing such as Collard-Wexler and Loecker (2013) who study innovations in the U.S steel industry. However, measuring gains from reallocation is more complicated in the retail sector. Unlike manufacturing, retail productivity at the establishment level is not readily available in the data as physical units of output are not defined in retail. Moreover, geography matters. Establishments compete in many local markets rather a single aggregate market which requires one to model substitution patterns by consumers.

1.1.3 Spatial competition

The structural model that is presented in Section 3 borrows heavily from the literature in industrial organization on spatial competition which feature models of discrete choice with a geography that reflects real world locations (Davis, 2006; Chiou, 2009; Holmes, 2011). Spatial competition is generated by assuming that customers face costs of traveling and hence prefer to shop at stores closer by, leading to a gravity based demand structure. A key difference here is that the model here is one of industry equilibrium. Firms and stores alter their behavior in response to changes in their operating environment. This is particularly useful when performing counterfactuals, which require equilibrium responses.

2 Data

The bulk of the data used in the paper come from the Census of Retail Trade (CRT) in the years 1997, 2002 and 2007. Although the CRT is available for every year ending in "2" or "7" dating back to 1977, I restrict the sample to 1997-2007 since online sales are virtually non existent before 1997 and the NAICS classification in the dataset is most stable across these years. Data from the CRT are derived from responses to forms mailed out to retail stores and administrative records. All establishments of multi-unit retail firms (chain stores), single unit establishments with approximately more than 3 employees and a sample of small employers receive forms requesting information. Information for stores not selected in this sample or not responding to mailed forms are derived from administrative records from other federal agencies including the IRS. Roughly 10 per cent of records are based on administrative records. The combined data represent the universe of retail stores (NAICS: 44-45 and 72) that exist in the US in a given Census year.

Data is at the store level and includes information on employment at March, annual sales, 6 digit NAICS code, geographic location (by zip code) and longitudinal firm and store identifiers. As the focus is on retail stores that can potentially be affected by online retailers, I restrict the data to those in retail industries that are traditional inventory-holding type stores. This excludes the following industries: accommodation and food services, non-store retailers, gasoline stations, automotive dealers, art dealers and mobile home dealers.

The data used to construct measures of internet access come from the Current Population Survey (CPS) Internet and Computer Use supplements in Dec. 1998, Oct. 2003, Oct. 2007 and Oct. 2012 which I will map to internet access in 1997, 2002, 2007 and 2012 respectively. Each survey asks a variant of whether anyone within a household has access to internet at home. I describe the estimation of zip code level internet access from individual level estimates in the Appendix.

Data on commuting distances and times were collated from Google Maps in June 2013. Distances reflect the shortest path by road network, including highways, between an origin and destination pair. Travel times reflect a combination of factors including posted speed limits and average traffic levels. In all, the sample of data I use has distance and travel times for 3.2 million unique origin-destination combinations in the US.

3 Model

The model that follows is centered on a discrete choice framework with consumers choosing where to purchase some real unit of consumption. Consumers receive utility from making purchases with preferences dictating the utility attached to each purchase option. Taking consumers' preferences as given, stores set prices to maximize profits. The model is estimated at the 6 digit industry which implies that parameters are industry specific and agents' decisions are within industry. For instance, the utility of purchasing online is industry specific and stores only worry about competition from other stores within the same industry. Book stores are only concerned about competition from other book stores.

Stores and households exist in their real world locations with households able to shop anywhere. In practical terms, a store or household's location is the center of the zip code

tabulation area in which they reside.[2] Geography matters since consumers face a cost of traveling to stores. From this assumption, two patterns will arise. First, all else being equal, consumers will prefer to shop at nearby stores. Second, the effective set of competitors for any store are other nearby stores. Therefore, a store's market power is largely determined by its location. Being surrounded by many other stores will lower market power for a store lessening optimal markups.

The online retailer that competes with brick and mortar stores is exogenous with characteristics that are inferred from Amazon.com annual reports.[3] The utility that consumers get from shopping online is a parameter that is estimated for each year and each industry while cost and markup information is inferred from annual reports. Treating the online retailer as exogenous is a deliberate and conservative approach that makes the fewest assumptions regarding Amazon.com's behavior.

One limitation of the model is that it does not allow for cross industry substitution. For example, the choice set of consumers deciding to purchase from bookstores are limited to bookstores alone which does not explicitly allow for substitution towards, say, general merchandise stores. While it would be straightforward to extend the model to all industries and to allow for rich substitution patterns via consumer-industry interactions, it is simply not computationally tractable given the size of my dataset. To do so, it would be necessary to reduce the number of parameters that reflect firm and store level heterogeneity so that the model is closer to a model of industry rather than store choice. Instead, I lean towards estimating a simpler model incorporating within industry choices to make use of the rich store level data.

3.1 Consumers

Consumers need to make multiple purchases of some real unit of consumption throughout a year. The real unit of consumption is fixed across stores in a given industry and represents a bundle of goods. For example, the real unit of consumption for grocery stores should be interpreted as a bundle of groceries. Consumers choice set includes each brick and mortar store, the online retailer (if they have access to the internet) or the option to undertake home production. Once they have chosen where to make their purchase, consumers can only purchase a single unit of consumption. This abstracts from intensive margin adjustment in the size of purchases at a store, although the decision to move from home production to a store/online retailer is akin to intensive margin adjustment.

Each consumer has preferences among these options which describe the utility conferred from purchasing at each option. Because the real unit of consumption is fixed across the options, differences in utility implicitly reflect differences in prices and the level of services, broadly defined, between options. A supermarket may be preferred to others because prices are lower and the customer service, opening hours and quality of produce is better.[4]

[2] Zip code tabulation areas represent contiguous geographic areas whereas zip codes are defined by the USPS and are collections of addresses.

[3] Amazon.com accounts for a sizable chunk of online retail sales, 11.2 percent of total retail e-commerce sales in 2011. Compare this to Walmart which accounts for only 5.4 percent of total retail sales in 2013.

[4] This assumption regarding the same real unit of consumption across all stores is necessary as the data only reveals nominal sales at stores. It is impossible to identify differences in real purchase sizes from the data.

Denote the number of purchases per person in location l as $m_{l,t}$ which for estimation purposes is a function of time and demographic characteristics in the location.

$$m_{l,t} = \lambda_t e^{\lambda_i hhinc_{i,t} + \lambda_f family_{i,t} + \lambda_e educ_{i,t} + \lambda_w white_{i,t}} \tag{1}$$

A fraction of customers in each location, $\nu_{l,t}$, have access to the internet and hence access to online shopping while the remaining fraction do not.

Prior to shopping, consumers receive idiosyncratic taste shocks for each purchase option and choose the option that gives them the highest utility. Choosing home production delivers utility u_h normalized to zero before the taste shock while shopping online delivers utility equal to $u_{o,t}$ if the customer has access to the internet. Denote the utility or consumer surplus that a consumer at l receives from shopping at store i as $u_{l,i}$.

$$u_{l,i} = \theta_{f(i)} + \mathbf{x_i}'\theta - c^d_{l,i} - \eta p_i + \xi_i + \epsilon_{l,i} \tag{2}$$

Because it is necessary to estimate $u_{l,i}$, I assume that it is comprised of some observable factors: a firm fixed effect, $\theta_{f(i)}$, fixed effects relating to unobservable store characteristics, $\mathbf{x_i}$, disutility from distance, $c^d_{l,i}$, disutility from prices, p_i, store fixed effect, ξ_i and the taste shock, $\epsilon_{l,i}$ which abusing notation, is specific to individuals.

The firm fixed effect is shared by all stores owned by the firm in the same NAICS category. It captures the common element of a firm's stores such as effects that are generated by the firm's brand and marketing. The store fixed effect captures effects specific to the store but not captured by other observables. As in Holmes (2011), store age is included as the store characteristic. I also include economic density, measured as establishments per square mile, as a store characteristic which captures utility from being nearby workplaces and other stores.

The disutility from traveling between locations consists of two components: a component that depends directly on the distance between locations and another component that depends on the opportunity cost of time, which is captured by the time taken to travel between locations interacted with household income (or household income divided by speed). Distances and travel times are measured by data on the distance and travel times by road between two zip code centoids from Google Maps. To control for heterogeneity in the size of zip codes, I set the within zip code travel distance to half of the nearest zip code.

$$c^d_{l,i} = \left(\xi_0 + \xi_1 \frac{\log\left(hhincome\right)}{speed_{l,i}}\right) d_{l,i} \tag{3}$$

The final term is the taste shock which each customer draws before choosing where to shop. I assume that the support of the shock is unbounded, implying that all stores face a positive probability of being chosen by a customer, no matter how poor their offering to customers. This reflects the randomness in reality that drives people to sometimes visit stores that are much worse on average than others.

Although the model is simple at this stage, I further simply by assuming that the taste shock is distributed i.i.d extreme value type 1. This generates an analytic expression for the probability that a customer from a certain location shops at a given store. With a smaller dataset or one based on an aggregate market, it is computationally feasible to assume a more general form for the error term and use simulation based methods to derive market shares to avoid the well known problems of using logit errors. However, this is not possible here as I

estimate the model on the universe of retail stores in the US, which makes simulating market shares for each store computationally intractable. Using the properties of the extreme value type 1 distribution and the independence assumption of the error terms, I can express the probability that customer j with internet access shops at store i as $s^j_{l,i}$ with internet access

$$s^j_{l,i} = \frac{e^{u_{l,i}}}{e^{u_h} + \mathbb{I}_j e^{u_{o,t}} + \sum_j e^{u_{l,j}}} \tag{4}$$

where $\mathbb{I}_j$ is an indicator for whether customer j has internet access or not.

Equation (4) shows the rival nature of competition. Any improvement in the utility that a particular store delivers to customers increases that stores share of customers while decreasing everyone else's. One property of this demand structure is that entry by a competitor leads to reallocation from all retailers, not only stores offering low utility due to the random component of utility. For example, a McDonalds that opens offering worse utility than all incumbent restaurants will take market share from all incumbents as the entry of McDonalds endows households with another opportunity for a large taste shock.

3.2 Retail stores

Given household preferences, each store sets a price that applies to all consumers in order to maximize profits. I assume that stores face a constant marginal cost of serving a consumer as well as a fixed operating cost. Denote π_i as gross profits for store i.

$$\pi_i = \max_p (p - c) y_i(p) \tag{5}$$

Total purchases at a store, y_i, are equal to purchases aggregated across locations.

$$y_i = \sum_l y^{access}_{l,i} + y^{noaccess}_{l,i} \tag{6}$$

Purchases from a type of consumer from a particular location are equal to the store's share of purchases multiplied by the total number of purchases in that location from that particular type of customer. Prices affect the purchases from consumers through the effect on the share of consumers or allocation across stores.

$$y^{access}_{l,i} = m_l n_l \nu_l s^{access}_{l,i} \tag{7}$$

$$y^{noaccess}_{l,i} = m_l n_l (1 - \nu_{l,t}) s^{noaccess}_{l,i} \tag{8}$$

The optimal price balances the extra revenue gained per purchase against the foregone profits from consumers that choose to shop elsewhere and satisfies the following markup rule over marginal costs.

$$p_i = c + \frac{1}{\eta\,(1 - \bar{s}_i)} \tag{9}$$

where $\bar{v}_i$ represents a store's purchases-weighted average market share across consumers with and without access to the internet.

$$\bar{v}_i = \sum_l \frac{y_{l,i}^{access} s_{l,i}^{access} + y_{l,i}^{noaccess} s_{l,i}^{noaccess}}{y_i}$$

Equation (9) shows that markups are driven by a common component and a component specific to stores. The common component is the extent to which customers dislike higher prices, captured by the parameter η. More sensitivity to higher prices will lead to lower markups. The variation in markups come from variation in stores' average market shares. Stores with large market shares have a lower own-price elasticity and hence set higher markups as the tradeoff between increasing margins and acquiring more customers tilts in favor of customer acquisition. Entry by a new establishment in an area lowers incumbents' market shares putting downward pressure on markups.

Average market shares can differ across markets for two primary reasons: market size and quality of retailers. An increase in market size, assuming that the number of stores is roughly proportional to market size, will lead to decreases in market shares. Each store will serve the same number of consumers but their market share is smaller. Hence the model predicts that markups are lower in larger markets, all else being equal. The quality of retailers also matters since customers can always choose the outside or online option. Competition amongst brick and mortar retailers in two areas may appear equal (in terms of each store's share of brick and mortar sales), yet if average quality differs between the two groups, actual market shares (measured against total sales including both home production and online sales) will differ. As an example, suppose a mediocre restaurant is the sole restaurant within 100 miles. Although that restaurant may appear to have a monopoly and hence unlimited pricing power, the fact that customers can cook at home and view the restaurant relative to this outside option implies that the restaurant's market power is in fact quite limited.

The final piece of the model concerns establishment exit. For simplicity, each firm is assumed to have a fixed cost of operating a store in a location which represents the fixed cost of labor, rent or opportunity costs of holding property and other implicit costs such as an entrepreneur's foregone wages or shareholder's required return on assets. I assume that firms own stores and are responsible for the entry and exit decisions of stores. Denote $V(i, \Omega)$ as the value of firm i which encompasses net profits at each of its stores given the aggregate state encapsulated in Ω.

$$V(i, \Omega) = \max_{S_i} \sum_{j \in S_i} \pi_j(S_i; \Omega) - \varphi_j + \beta \mathbb{E} V(i', \Omega') \tag{10}$$

where S_i is the set of stores operated by firm i and $\pi_j(S_i; \Omega)$ are the gross profits from store j making explicit their dependence of the aggregate state and the set of stores operated by the firm. Optimality regarding the extensive margin of store choice therefore requires that perturbations to the set of stores to be suboptimal. That is, any store expansions or closures in the current period must not generate additional net profits to the *firm*. This rules out new stores that may be profitable on their own but not to the firm due to cannibalization effects. More concretely, $\forall i, \Omega, S'$,

$$V(i, \Omega) \geq V(i, \Omega; S') \tag{11}$$

where $V(i, \Omega; S')$ is the firm's value with the set of stores S'.

3.3 Aggregate welfare

Aggregate welfare is the total surplus (consumer surplus and profits) generated from all purchases net of the operating costs of stores.

$$W = \sum_l \sum_i y_{l,i} \left(\frac{u_{l,i}}{\eta} + p_i - c \right) - \sum_i \varphi_i \tag{12}$$

Note that consumer surplus from a purchase is divided by η which scales utility in dollar terms implying that prices have no bearing on the surplus at a store, beyond the allocation of households across stores. Equation (12) shows that what matters for welfare in the model are the services that are provided in the process of consumption. Retail stores generate welfare by providing services to consumers that have value over and above the value from the goods purchased. The value of these services are large in reality, with value added from Retail and Wholesale Trade accounting for 51 per cent of personal consumption expenditures on goods on average from 2000-2012.

Normalizing aggregate welfare by the number of purchases generates a new measure of retail productivity. This new measure captures an economy's ability to allocate consumers efficiently such that in the process of consumption, a relatively valuable amount of services are generated at relatively little cost. Other measures of retail productivity such as value added or sales per worker get at a similar notion except that they suffer from issues due to the nominal measure of output. A social planner using these standard measures of retail productivity would prefer low cost monopolies charging high prices. Instead, using the measure in this paper would lead the social planner to prefer low cost firms that provide relatively valuable services.

3.4 Equilibrium

Equilibrium in the model is optimality by consumers, which is satisfied by equation (4), optimality by stores, which entail a set of prices for every store that satisfies equation (9) and extensive margin choices by firms which satisfy equation (11) and market clearing. Regardless of prices, goods market clearing in each location is satisfied since aggregate demand is fixed with any residual demand not satisfied by stores going to either the outside option or an online retailer. Another way to think of this is that rather than carrying inventory, stores can instantaneously produce the good and supply it to the consumer. Market clearing holds since any pattern of customer choices across stores, home production and online retail are instantly fulfilled by that option, guaranteeing market clearing.

Theorem 1. *There exists a unique equilibrium to the model.*

Proof. See Appendix. □

Because the primary goal of the model is to Improvements by online retailers increase the utility from purchasing online and have two effects. The first of these channels is straightforward. Holding fixed allocations, consumer surplus increases for those already shopping

online leading to aggregate welfare gains. Second, the market structure of retail changes as consumers substitute away from brick and mortar stores towards online retail. Sales at brick and mortar stores fall in addition to markups as stores' market power falls. Both of these lead to store closures which has an ambiguous effect on aggregate welfare. Consumers who initially substitute to online retail clearly gain but those consumers who still preferred to shop at the subsequently close store are now forced to find the next best alternative. Offsetting some of these potential losses are the savings in the closed store's fixed operating costs.

4 Estimation

The estimation procedure follows two steps. First, I estimate the demand system which uncovers parameters that govern consumer surplus on offer at each store and the online retailer. This requires estimates of internet access by zip code, demographic data by location, distance and travel times between zip codes, data on store characteristics and local market sales shares. When estimating the demand system, I do not need estimates of fixed costs since I merely need to assume that fixed costs are such that they rationalize the network of stores observed in the data. In the final stage of estimation, I estimate bounds for each store's fixed operating costs using a revealed preference approach based off the estimated demand system.

4.1 Demand system

Each store's consumer surplus is comprised of a firm fixed effect, store fixed effect, store characteristics, disutility from prices and the cost of distance from consumers. The basic strategy is to match each store's observed share of sales within its zip code to that generated by the model. A store observed with a high market share will have its consumer surplus parameters tweaked until the model generated market share matches the data.

More precisely, each parameter is identified by assessing the impact of variation in the variable of choice on stores' market shares. For instance, the parameter associated with store age is identified by the change in market share caused by variation in store age. However, note that since market shares are calculated within zip codes, this strategy only identifies parameters that rely on within zip code variation. To identify parameters associated with variables that vary only across zip codes, I use variation in the level of sales. For example, distance from households only varies for stores in different zip codes. The cost of distance is inferred from the effect on zip-code aggregated sales from being located in different areas, perhaps for instance, further away from a major population center.

The remaining component of stores' consumer surplus that needs to be identified are the store fixed effects. Here I follow Berry et al. (1995) in setting store fixed effects such that together with the observable component of store utilities, the model generates market shares equal to the observed market shares in the data. These residual store fixed effects account for the portion of sales not accounted for by observable characteristics. Note that finding a set of residual store fixed effects that match observed market shares is a fixed point problem as changing a single store fixed effect alters all market shares. Moreover, this task is made more difficult than the standard Berry et al. (1995) algorithm as for any set of store fixed effects, a fixed point for prices needs to be found.

Formally, I want estimates of fixed effects and prices for each store such that (i) stores estimated shares of sales among stores located in the same zip code are equal to that in the data and (ii) the pricing equation holds.[5] Note that this is a non trivial problem because solving for equilibrium prices is itself a fixed point problem. While it appears that some form of a nested fixed point algorithm is necessary to find a solution, it is possible to find a solution that iterates on a transformation of the store fixed effects alone. I provide a more formal description of the algorithm in the Appendix and a proof that the operator defined by the algorithm is a contraction and provides optimal prices and store fixed effects that match observed market shares.

Because the data covers the universe of retail stores in the US, I make some restrictions in estimating the model to make the estimation procedure computationally tractable. I reduce the parameters that the estimation procedure needs to evaluate by only estimating firm fixed effects for firms with greater than 10 unique stores in a 6 digit NAICS industry in a single year over the sample. For the remaining small firms, I estimate a fixed effect term for each separate year, which reflects the average utility a small firm delivers relative to the outside option in a given year. To compensate for the loss in precision for these smaller firms, I include firm size as an observable characteristic.

4.1.1 Instruments and identification

The demand system is estimated via two-step GMM which requires as many moments as parameters. There are three broad classes of moments that I use. The first are moments based on the assumption that the residual store fixed effects, ξ_j, is mean independent of the variables that vary across stores within a location, $\mathbf{z}_1$.

$$\mathbb{E}(\xi_j|\mathbf{z}_1) = \mathbb{E}(\xi_j) \tag{13}$$

The residual store fixed effects reallocate market shares within locations to match observed market shares. Hence, these moments can only identify parameters that relate to variables that are capable of changing market shares within locations. The moments that I use in estimation require that the covariance between the unobservable store term and the variable k is zero. The variables used with this moment include firm size by number of stores for small firms, an indicator variable equal to one if the store is more than a year old, an indicator that selects small firms' stores for each year and an indicator that selects all stores owned by a firm, for each firm.

$$\mathbb{E}_j(\xi_j z_{k,j}) - \mathbb{E}_j(\xi_j)\mathbb{E}_j(z_{k,j}) = 0 \tag{14}$$

Each of these moments identifies a single parameter. As an example, consider the moment that selects stores for a given firm. Suppose that the moment were positive for a given parameter vector. This implies that stores owned by the firm require a higher than average store fixed effect to match market shares. To move the moment closer towards zero, it is necessary to increase the firm's fixed effect parameter to soak up the positive covariance.

The next set of moments are based on measurement error between the level of sales in an area and the model's estimate of an area's sales. Measurement error still exists even when

[5] Because there are multiple years in the data, I choose to match market shares aggregated over the 3 separate years in the data. While it is possible to match market shares in each year, this raises the risk of overfitting.

market shares are equal to that in the data because market shares are calculated within, not across locations. The moments based on measurement error are based on the assumption that the conditional mean of the measurement error with respect to the variables that vary across locations is zero.

$$\mathbb{E}_j(e_j|\mathbf{z}_2) = 0 \tag{15}$$

Moments based on measurement error are of the form

$$\mathbb{E}_j(e_j z_{k,j}) = 0 \tag{16}$$

The moments used are based on: distance to customers, distance to customers times population density and population weighted demographic variables. As an example, consider the moment based on the distance to customers. If the moment is negative, it implies that stores relatively close by to customers have larger measurement errors. To remove the correlation, it is necessary to increase the constant parameter relating to distance in the utility function which has the effect of reallocating customers and consequently sales towards stores relatively close by to customers. Similar arguments can be made for the other moments based on measurement error.

With this set of moments, the model remains under-identified. To complete the set of moments, I use moments based on more aggregated data. The demand shifters in each year are identified by yearly differences in total brick and mortar retail sales between the model and the data. Customers' disutility from prices are identified via the markup rule for prices and the log difference between average markups in the model and in the data. The data on average markups are calculated from the average markup between 1997 and 2007 for each industry from the Annual Retail Trade Survey.

The remaining estimated parameters in the model are the utilities from the online retailer. I use moments that state that the model's estimate of the online retailer's share of total sales is equal to that in the data. Bear in mind this parameter is 6 digit industry specific/ To determine the online market share for an industry, I calculate the online market share for each product and then use the weight of each product in an industry's sales to compute a weighted average of the online market shares across products. There is substantial heterogeneity in these market shares, with online market shares for book stores and meat markets 24 and less than 1 per cent respectively. The full list of moments and the parameters they each identify are listed in Table 1.

Finally, marginal costs for brick and mortar stores are normalized to 1 since the Census of Retail Trade is uninformative regarding costs at each store. Since prices follow a markup rule over marginal cost, normalizing costs shifts any cost-driven variation in prices between stores into the consumer surplus term as discounts. Marginal costs for the online retailer are set to operating expenses over the cost of goods sold from Amazon.com annual reports, relative to the same calculation for the entire brick and mortar retail sector from the Annual Retail Trade Survey.

Normalizing costs implicitly fixes the level of prices which is likely to distort the amount spent per trip in some industries. Instead, variation in the amount spent per trip is soaked up by the cost of distance parameter. Estimates of the cost of travel will be lower in industries where consumers spend a large amount per trip. As a result, the model will suggest that consumers in these industries make many trips in order to make up the dollar volume of sales.

Table 1: List of moments

Moment	**Parameter identified**
Local market shares	
indicator for groups of stores by firm	β_f
indicator for stores of small firms by year	$\beta_{s,t}$
firm size (# stores)	θ_{fs}
dummy for store age ≥ 2	θ_{age}
Measurement error across locations	
household income	β_i
family share of households	β_f
education	β_e
white share of pop.	β_w
distance	ξ_0
distance $\times$ log household income / speed	ξ_1
Aggregate moments	
average markup	η
level of sales for each year	λ_t
online market shares for each year	$u_{o,t}$

Regardless, in aggregate, the total expenditure and total costs of travel are the same.

4.1.2 Demand system parameter estimates

While I estimate 17 industry-specific parameters for each of the 29 industries and firm fixed effects for every firm with at least one year of operation with more than 10 stores in some retail industry, I am prevented from reporting estimates at the firm, industry or year level due to disclosure restrictions imposed by the US Census Bureau. Instead, I report the mean of these parameters estimates and their standard errors across industries and across years where applicable. To help get a sense of the dispersion in these parameters and their standard errors across industries, I also report the 25th and 75th percentiles of the parameter estimate and the respective standard error.

Most of the parameters are relatively well identified since the values over which the standard errors range are relatively small compared to the parameter estimates. The estimated parameters relating to firm size and age are consistent with the notion that firms with a greater number of stores offer greater utility to customers relative to smaller stores and that stores less than 2 years old have lower sales than older stores. To assess the validity of these parameter estimates, I calculate moments implied from the model that were not targeted in the estimation process and compare them with data where possible. These include moments regarding shopping related travel of consumers, the variance of prices and the relationship between markups and competition.

Table 3 shows the effect of distance on the demand of an arbitrary store using the mean of estimated parameters and in a location with the average driving speed of Los Angeles, 26.8 miles per hour.[6] The table calculates the demand of a store located a certain distance

[6]This is calculated by taking the mean of all unique zip code pairs with one side of the pair in Los Angeles.

Table 2: Summary of parameter estimates and standard errors

Parameter	Estimate (std. err.)[1] Mean	25th	75th
Utility parameters			
online shopping constant - u_o	-2.25 (0.38)	-4.21 (0.80)	-0.55 (0.03)
distance constant disutility - ξ_0	0.120 (0.004)	0.058 (0.0020)	0.187 (0.007)
distance opportunity cost - ξ_1	0.007 (0.001)	0.001 (0.0001)	0.013 (0.001)
disutility of prices - η	2.34 (0.28)	1.29 (0.15)	2.62 (0.32)
firm size - θ_{fs}	0.51 (0.11)	0.27 (0.06)	0.66 (0.12)
age dummy - θ_{age}	0.40 (0.07)	0.28 (0.05)	0.54 (0.09)
economic density - ϕ	0.046 (0.003)	-0.19 (0.02)	0.042 (0.005)
Demand parameters			
demand ('000 visits per year) - λ_s	0.30 (0.05)	0.001 (0.0002)	0.13 (0.02)
household income demand - β_i	0.38 (0.06)	-0.44 (0.11)	1.36 (0.22)
family share demand - β_f	0.06 (0.01)	-1.22 (0.17)	1.69 (0.25)
white demand - β_w	1.10 (0.15)	0.41 (0.03)	1.69 (0.25)
education - β_e	0.08 (0.05)	- 0.90 (0.14)	1.49 (0.31)

[1] Percentiles represent the parameter estimate and standard error for the industry representing the percentile.

away from a customer relative to the same store located next door to a customer. Demand falls both with distance from a customer and household income. Holding household income at $25,000, demand falls by 25 and 76 per cent as distance increases by 1 and 5 miles. In an area with $100,000 in household income, demand falls slightly more over those distances, 26 and 78 per cent respectively as opportunity costs are larger.

Table 3: Relative demand to a store 0 miles away

Miles from customer	Household income $25,000	$50,000	$100,000
1	0.75	0.74	0.74
2	0.57	0.55	0.54
3	0.43	0.41	0.40
4	0.32	0.31	0.29
5	0.24	0.23	0.22
10	0.06	0.05	0.05
25	0.00	0.00	0.00

[a] Calculated at 26.8 miles per hour average travel speed, average speed for Los Angeles.

Table 4 shows various moments regarding travel behavior of consumers in the model, averaged across industries. The median distance travelled per year is estimated to be 359.8 miles with a median cost of travel of 14 cents per mile with roughly 5 cents per mile of that cost coming purely from the cost of driving. Five cents per mile is lower but comparable to 14.9 cents per mile, the estimate of variable per mile driving costs provided by AAA for a medium sedan in 2007. The model's estimate of average distance per trip is also slightly higher than in the data. The average shopping trip in the Department of Transportation's

National Household Travel Survey 2009 was between 7.2 to 7.63 miles whereas the median trip was 8.05 miles in the model.[7]

Table 4: Implied dollar costs of traveling to stores

	Percentiles across population weighted zip codes[1]						
	5	**10**	**25**	**50**	**75**	**90**	**95**
Median across industries							
Avg distance per trip (miles)	3.41	4.29	5.91	8.05	12.38	16.8	19.95
Yearly distance (miles)	30.2	66.3	230.0	359.8	518.8	729.9	899.4
Cost per mile ($)	0.12	0.13	0.14	0.14	0.15	0.16	0.16

[1] Percentiles are calculated for each industry and then averaged across industries.

In the model, variation in prices comes purely though variation in competition across geographic space. To show that the estimated model is quite sensible with regards to price setting, I benchmark the model against a comparable estimate of the variance in prices across a geographic area. The median standard deviation in log prices, shown in Table 5, is 2.6 per cent which is slightly lower than the estimates of 3.3 to 5 per cent estimated in Kaplan and Menzio (2013) for the variance in prices within MSAs due to slow moving store characteristics such as store quality and location. Kaplan and Menzio (2013) interpret the variance in prices as evidence of search frictions, whereas the results here suggest that a large part of the variance can be explained by variation in competition within MSAs.

Table 5: Variation in prices across markets

	Percentiles across unweighted zip codes[1]						
	5	**10**	**25**	**50**	**75**	**90**	**95**
Std. dev (log prices)	0.000	0.001	0.011	0.026	0.045	0.074	0.101

[1] Percentiles are calculated for each industry and then averaged across industries.

To assess the relationship between the *level* of markups and demographics, I regress the average log markup paid in a 50 mile area on demographic variables describing the market and report the results in Table 6. Bear in mind that these are model estimates of markups, not actual markups from the data. Regardless, the model suggests that markups fall by population density, household income and education level across a broad range of industries. The effect of the non-latino white share varies across industries. A doubling of the population density, household income and the education level of an area is associated with an average 5.5, 4.4 and 2.7 per cent fall in markups. This is consistent with empirical evidence that competition is stronger in markets with greater population density (Bresnahan and Reiss, 1991; Campbell and Hopenhayn, 2005).

4.1.3 Estimates of store-level productivities

Compared to a standard heterogeneous firms model, a key difference is that the distribution of store surpluses (or)productivities varies across space. The mix of stores does not repeat

[7] I restrict the sample to only include trips undertaken by automobile and from stores to home or vice versa.

Table 6: Markups and demographics

	Regression[1]		
	Avg	**25th**	**75th**
Log population density	-0.055	-0.091	-0.006
Log household income	-0.044	-0.058	-0.011
Log frac. pop non latino white	0.026	-0.013	0.081
Log frac. pop > bachelors degree	-0.027	-0.064	-0.006

[a] All coefficients are significant at the 1% level.

itself across geographic space so that the distribution of stores in each market is identical to one another. Some areas are populated with higher quality flagship stores of retail firms while others have less well maintained stores. Since estimates of these store level surpluses are available, it is possible to assess whether the distribution of store surpluses varies systematically. Define a store's surplus per purchase as the consumer surplus in dollar terms plus profits, per purchase by a household.

$$z_{l,i} = \frac{u_{l,i}}{\eta} + p_i - c \tag{17}$$

Prices have no effect on a store's surplus since prices enter negatively in consumer surplus and only serve to split the surplus between consumers and the store. To detect systematic variation in the distribution of store-level surpluses across space, I relate the mean store surplus in an area to demographic variables in an area. Table 7 describes the correlation between the mean store surplus and demographics at the zip code level and results from regressing the mean store surplus in an area on demographic variables. To facilitate comparisons between industries, the measure of mean store-level surplus in the regressions is demeaned at the industry level.

Table 7: Mean store-level surplus and demographics

	Correlation			Regression[a]		
	Avg	**20th**	**80th**	**Avg**	**20th**	**80th**
Constant	-	-	-	0.13	-0.86	0.84
Log population density	0.41	0.30	0.56	0.25	0.06	0.44
Log household income	0.06	-0.00	0.12	-0.12	-0.17	0.05
Log frac. pop non latino white	-0.16	-0.22	-0.10	-0.05	-0.09	-0.01
Log frac. pop > bachelors degree	0.24	0.14	0.33	0.21	0.05	0.35

[a] All coefficients are significant at the 1% level.

The results show a strong correlation between the mean store surplus and both the education level and population density in an area. Correlations for both population density and education are positive, even in the 20th percentile industry, suggesting a robust effect across industries. A doubling of population density and education is associated with an increase in the mean store surplus of roughly 25 cents per purchase (6.7 per cent of the nationwide average store surplus) and 21 cents per purchase (5.6 per cent). In contrast, household income and the fraction of white households have an ambiguous relationship with the mean store

surplus in an area. This perhaps reflects the industry-specific relationship between demand and these other demographic variables.

4.2 Fixed costs

For a set of estimated parameters, the model generates corresponding estimates of profits for each store in equilibrium. Hence, by varying parameters which reflect changes in the operating environment, one can measure the equilibrium effect on the retail sector. In calculating these responses, it is necessary to obtain estimates of fixed costs as establishment exit is likely to be important when considering the effect on online retail.[8]

One way to measure fix costs is directly from the data, as Holmes (2011) does for Walmart. Unfortunately, data on costs by store in the Census of Retail Trade is restricted to payroll data which is a noisy measure of fixed cost. A large component of payroll probably reflects variable costs and fails to account for other important fixed costs such as rent or the opportunity cost of funds. Another approach would be to estimate fixed costs that are specific to locations using entry and exit patterns for all stores location by location. Roughly speaking, the estimated fixed cost in a particular location would be the threshold level of profits that induces exit. However, this fails to account for the massive heterogeneity that is likely to exist between stores in their fixed costs. Smaller firms probably have much lower fixed costs per store than large national retailers. A model estimated in this way would overstate the role of exit since low cost/high surplus stores would exit at too fast a rate relative to the data and vice versa for high cost/low surplus stores.

Instead I use a revealed preference approach that identifies store-specific fixed costs through the geography of a firm's network of stores. The main disadvantage of this approach is that it only set identifies fixed costs. Hence I will report results from counterfactuals from the midpoint these sets, with the full set of results in the Appendix. Results computed from the lower bounds of fixed costs are conservative estimates of exit, since it requires a relatively large reduction in sales to induce exit of any store. On the other hand, the results using the upper bounds are less conservative as exit thresholds are passed more easily.

To construct these bounds, I exploit conditions that describe the optimality of each firm's choice of the network of stores to operate. Informally, fixed costs must be low enough that it would not be optimal for firms to close a store in a location (where a store already exists) and high enough that firms do not find it optimal to open another store in the same location. These perturbations are similar in vein to the pairwise deviations in Holmes (2011) and the one-step deviations in Morales et al. (2013) which use minor deviations in optimal policies to generate inequalities that set identify parameters. To identify a firm's lower bound for the fixed cost in a particular location, define a perturbed set of stores from the firm's optimal set of stores by adding an additional store to that location. Denote the optimal set S_i and the perturbed set $S_i(+, l)$. It must be the case that since we observe the firm operating S_i stores

[8]Entry is also likely to play a role but it is difficult to introduce entry into a model that covers the geography of the entire US and has a large number of potential entrants. Exit is less troublesome since it involves stores that exist and hence are observed.

in the location, operating $S_i(+,l)$ stores is suboptimal.

$$\begin{aligned} V(i;S_i) &\geq V(i;S_i(+,l)) \\ \varphi_{l,i} &\geq \pi_l(S_i(+,l)) - \pi_l(S_i) \end{aligned} \tag{18}$$

Similarly, denote the perturbation $S_i(-,l)$ which is identical to the optimal set of stores except that a single store is closed in location l. Optimality requires the following hold.

$$\begin{aligned} V(i;S_i) &\geq V(i;S_i(-,l)) \\ \varphi_{l,i} &\leq \pi_l(S_i) - \pi_l(S_i(-,l)) \end{aligned} \tag{19}$$

Equations (18) and (19) define the set of feasible estimates for firm-location specific fixed costs. The lower bound at a location is the increase in gross profit from adding another store in that location, accounting for the effect on price setting by all stores owned by the firm.[9] Another way to interpret these sets are as discrete analogs to the first order conditions that would come from optimization if the number of stores in each location were continuous rather than restricted to integers. Firm-location fixed costs would be identified from a firm's marginal profit with respect to the number of stores at that location if the number of stores were continuous. Because store choice is not continuous, it is not possible to use infinitesimal variations in store policy and hence only upper and lower bounds can be derived.

These bounds reflect the curvature of the firm's gross profits in a location with respect to the number of stores with the size of the gap between the upper and lower bounds dependent on the concavity of gross profits. Gross profits are concave both because of the structure of the logit probabilities and through cannibalization of profits at stores in other areas. Data on ownership are derived from the Report of Organization which has comprehensive data on the organizational structure of multi-unit firms. Every establishment in the Census of Retail Trade is effectively linked to the highest level firm that has a controlling interest (or chain of controlling interests) in the establishment. Therefore, cannibalization effects incorporate the reduced sales from all stores owned by the firm, not only stores with the same trade name. Failing to account for cannibalization effects for the full set of stores leads to estimates of marginal profits that are too high, biasing both bounds upwards and estimates of fixed costs that are too high, overstating the role of exit.

An advantage of this approach is that the bounds of fixed costs incorporate all factors that firms consider as fixed costs to operating a store, even unmeasured fixed costs. Even if one had complete accounting data for each store, unmeasured costs such as required return on equity vary across firms and affect the exit margins for firms. These unmeasured costs are accounted for here since they are inferred from a firm's choice of the number of stores to operate in a given location. One may be concerned that these estimates of the bounds for fixed costs encompass costs that are not traditionally associated with the fixed cost of operating a store, such as additional financing costs imposed on smaller firms or an owner's personal disutility (in dollar terms) of operating an additional store. However, it is precisely these costs that make this approach more useful since these costs are likely to influence firms' behavior and are unable to be quantified. A firm may decide not to open an additional store because it's unmeasured fixed cost is very large, even though its profit net of measured fixed

[9] I assume that firms hold prices at stores owned by other firms fixed.

cost for its sole store may suggest that expansion would be viable.

To understand the estimates of fixed costs that are inferred from store location patterns, it is useful to describe fixed costs from various patterns of store location and geography, assuming only for simplicity that we observe stores operating at max one store in a location. For a given location, larger fixed costs are inferred from stores generating higher gross profits relative to other stores in the same location, holding constant cannibalization effects. If fixed costs were not higher for these stores, it would otherwise be optimal for the corresponding firm to expand in that location. Hence it is not necessarily the smallest stores in an area that are the most sensitive to increases in competition from entrants. For single unit firms (mom and pops), net profit margins are smaller in markets with larger numbers of stores since there is less curvature in gross profits. Marginal gross profits are flat since the market is already saturated with stores, lessening cannibalization effects.

4.2.1 Fixed costs estimates

The estimates of fixed costs for each store are summarized in Table 8. The median estimate of the midpoint of fixed costs is $196,000 while the 10th and 90th percentiles are $35,000 and $1,062,000. To get a sense of how large fixed costs are relative to gross profits, I report the midpoint of fixed costs relative to gross profits. The 10th, 50th and 90th percentile estimates of fixed costs are 58.1, 94.9 and 99.8 per cent of gross profits respectively. The range of estimates is also relatively small, with the median difference between the upper and lower bounds of estimates for fixed costs representing 3.59 per cent of gross profits. I also report the percentiles of fixed costs relative to the median estimate in the same industry, as a gauge of heterogeneity in fixed costs within narrow retail industries. The heterogeneity within industries is substantial with the 10th and 90th percentile stores exhibiting fixed costs 0.34 and 3.65 times the median store in their industry, indicating that profitability matters for exit.

Table 8: Estimates of fixed costs

	Percentile						
	5	**10**	**25**	**50**	**75**	**90**	**95**
Fixed cost ($ '000s) - midpoint	16.1	33.4	96.1	249.5	590.2	1,254.5	1,968.4
Fixed cost (midpoint) - % of gross profits	45.3	64.1	84.3	95.0	98.8	99.8	99.9
Range of estimate - % of gross profits	0.01	0.03	1.55	5.64	16.14	34.21	48.70
Range of estimate ($ '000s)	0.03	0.21	2.47	16.81	80.79	297.73	647.62
Relative to median in industry	0.17	0.28	0.53	1	1.83	3.58	4.94

5 Results

To gauge the effect of online retail on the retail sector and the Marketplace Fairness Act (MFA), I use counterfactual exercises that hold other parts of the model constant. The first set of these counterfactuals are based on changes that are likely to have had a role in increasing online sales. These counterfactual exercises are like exogenous shocks which I impose on the model and then measure the resulting effects.

The MFA is assessed by introducing a tax on the online retail good which mimics the proposal to allow states to compel online retailers to collect sales taxes based on the destination of the shipped goods. I apply sales taxes to online retailers based on the state of residence of households and gauge equilibrium responses by both stores and customers. Effects of the rise in sales taxes depend in part on price elasticities, but also on local market conditions. In areas where operating margins are relatively tight, relatively small increases in taxes on online retailers can have a large impact on mitigating store closures and the subsequent reallocation that unfolds with store closures. On the other hand, effects in areas with fat operating margins depend primarily on price elasticities alone.

The counterfactual exercises are implemented on the 2007 data. Therefore, counterfactual exercises represent changes in online retail from 2007 to 2012 and effects from the MFA if it had been implemented in 2007. Since data from the 2012 Economic Census are unavailable, I compare the effects of the counterfactual exercises relative to the 2007 baseline as a means of assessing the magnitude of the effects. Although I can not directly validate the model's predictions since 2012 data are not unavailable, I show that the model is consistent with reduced form measures.

Computing these counterfactuals all rely on varying some element of the model and assessing responses by firms which requires computing equilibria. Equilibrium consist of a set of prices that satisfy store optimality taking other stores' prices as given and optimal exit decisions by firms. Computing the set of optimal prices is straightforward for a given set of stores as I show in the Appendix that one can write a fixed point algorithm for prices that is a contraction mapping and hence converges to a unique fixed point from any initial guess. The difficulty in computing equilibrium is in determining the set of stores that exist and satisfy firm-level optimality given the set of fixed costs for each firm. I describe the algorithm that finds an equilibrium in detail in the Appendix. Briefly, the algorithm is consistent with the "natural equilibrium" in Abbring et al. (2012), where multiple equilibria are resolved by selecting the equilibrium consistent with the weaker stores exiting first.[10]

5.1 Improvements in online retail

It is useful to frame the potential mechanisms that have lead to increases in the online share of sales in relation to the mail order industry, which offered a set of products via physical catalogs available to be shipped direct to customers. Superficially, online retail simply offers an improvement in this catalog technology with better features such as the ability to find products quickly, reviews to support decision making and a broader range of products available. These features can broadly be categorized as improvements in the services offered to customers above the value of the product itself and are unfortunately hard to measure quantitatively. However, the stages that bookend the online retail experience are measurable: access to the internet (catalogs) and shipping of goods to customers. I measure effects on welfare gains from these two measurable changes, and make up for the residual increase in the online share of sales by an increase that resembles online retailers' expansion into new industries.

[10] There are many other ways to implement equilibrium store closures. The other extreme is to close stores making the smallest losses first. However, experiments with algorithms that closed these relatively profitable stores first were unstable.

5.1.1 Increased internet access

Increased internet access expands the base of potential consumers for online retailers.[11] Figure 2 depicts the growth in internet access from 2007-2012 which I derive from CPS data on individuals' access to the internet.[12] While the mean increase in internet access across zip codes is low, the distribution is skewed with some zip codes experiencing a relatively large increase in internet access. To isolate the effect of the rise in internet access, I increase internet access for each location in the model by the estimated value from 2007 to 2012.

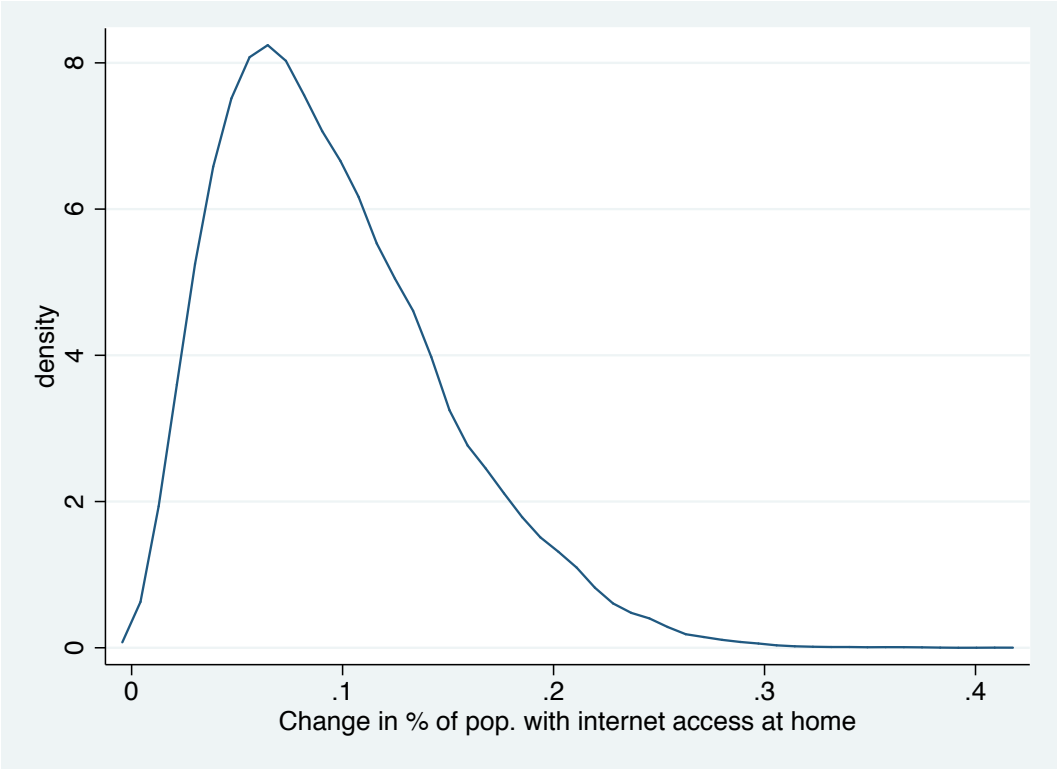

Figure 2: Estimated improvements in internet access across zip codes (2007 to 2012)

Improvements in internet access may have limited effects in terms of generating sales for online retailers because improvements are concentrated in areas with poorer households, who have less spending power. An industry that benefits from wealthier households is unlikely to experience an increase in market access to its desired customer base. On the other hand, improvements in access may lead to relatively large effects if retail stores serving these poorer areas are of relatively lower quality. If that is the case, customers in poorer areas will shift away from traditional retailers more so than customers in less poor areas, compensating for their relatively lower spending power.

[11] Greater internet penetration is similar to greater market penetration as in Arkolakis (2010) except that market penetration costs are paid by consumers.

[12] The derivation of zip code level estimates of internet access are provided in the Appendix.

5.1.2 Improved shipping practices

In contrast to increased internet access, there is no obvious way to introduce improvements in shipping into the model. Improvements in shipping are likely to be reflected in increased utility from purchasing online and changes in the end-to-end costs of delivery. The difficulty is in calibrating these changes so that they accurately reflect actual improvements in shipping.

The increased utility from improvements in shipping are based on the notion that delivery is valued since it eliminates the need to travel to stores. Instead of traveling, courier services such as UPS deliver goods from warehouses which in effect, shifts the commute cost to these courier services. Hence the value of delivery to consumers is the value of eliminating of travel to alternate stores, which varies across consumers. Consumers with access to relatively high quality stores nearby and low costs of commuting are inferred to value delivery less.

To determine the magnitude of the increase in utility, I begin with the assumption that the average disutility per trip experience by a consumer is equivalent to the value of same-day delivery, a shipping experience that replicates traveling to stores. Given this baseline, I calculate the increase in utility from 2007 to 2012 by determining how far away the industry is from same-day delivery in 2012 and measuring the progress made between 2007-12.

Delivery times can be reduced either by reducing the distance from the warehouse to customers or by increasing the speed at which goods travel which includes reducing bottlenecks in the supply chain. I use data on these margins to calculate reductions in shipping time. Specifically, I use the expansion of Amazon.com's shipping and fulfillment warehouses across geographic space over time and, productivity improvements at UPS and FedEx, who handle between 80-88 per cent of parcel deliveries. Given data on these margins, I calculate improvements in shipping by assuming that delivery times are proportional to the average distance between households and warehouses divided by productivity.

Figure 3 plots the expansion of Amazon.com's warehouses in terms of square feet and number of states. In 1997, Amazon.com operated warehouses in Delaware and Washington, one on each coast ostensibly to minimize transportation costs. By 2012, there were 14 states with warehouses spread almost uniformly over the US which reflects a large reduction in delivery times for customers. Note that this expansion is not simply maintaining warehouse capacity with the increase in sales since Amazon.com could expand warehouse space intensively in existing states.

To measure productivity, I calculate the total volume of packages delivered per real unit of expenses for UPS and FedEx where nominal expenses are deflated using the CPI.[13] Productivity gains are a reasonable measure of improvements experienced by end-users since margins at UPS and FedEx remain flat over the period. Figure 4 plots the measure of productivity from 1998 to 2012. From 1998 to 2012, the number of packages delivered per dollar of real expense increased 53.9 per cent with 32.8 percentage points of the improvement coming between 2007 to 2012.

Assuming that warehouses are uniformly distributed, the area covered by each warehouse fell from 0.986 to 0.211 million square miles or roughly 79 per cent between 1998 and 2012. Combined with the estimates of productivity improvements, I estimate that shipping times for the average household in 1997 were 7.14 times longer than in 2012 and 4.43 than in 2007.

[13] Using the PPI or implicit price deflator for the parcel delivery industry would result in an even slower increase in real expenses, implying that the estimates of productivity gains are conservative.

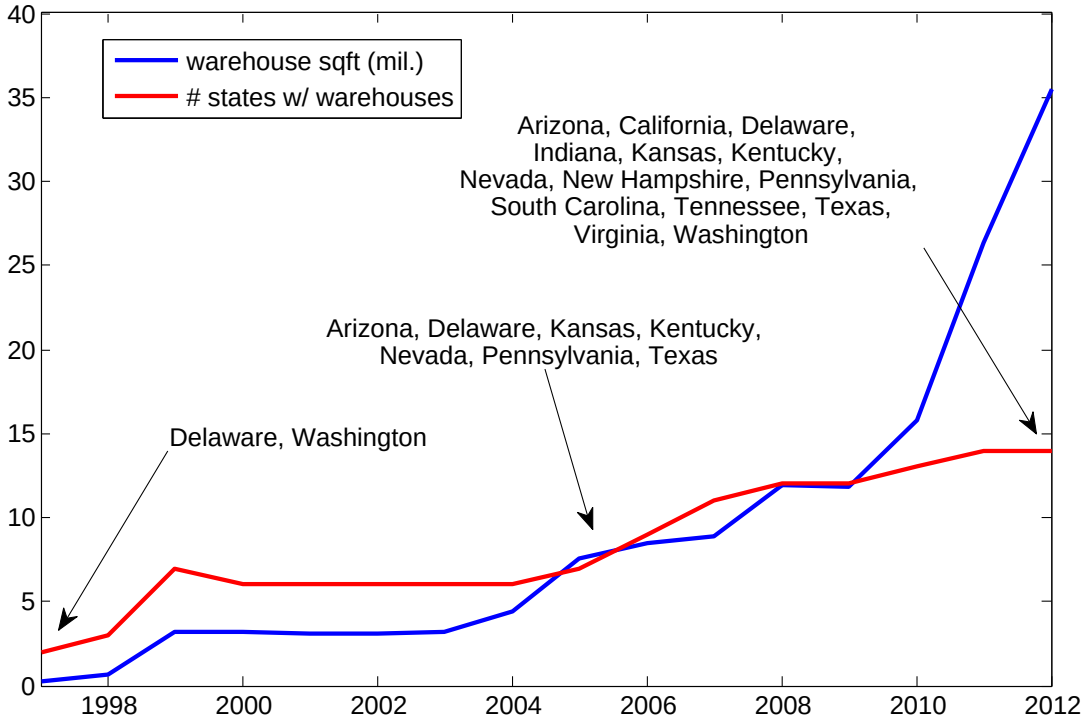

Figure 3: Shipping and fulfillment expansion at Amazon.com

The remaining step is to determine the the improvements necessary from 2012 to make same day shipping possible. Amazon.com's financial statements and annual reports suggest that same-day shipping will be available to the majority of the US population by the end of 2014.[14] By 2014, Amazon will have expanded its warehouses into Connecticut, Florida, New Jersey and Wisconsin, and productivity is estimated to rise by 7.1 per cent implying a 26.4 per cent reduction in shipping times from 2012. Therefore, I infer that a 55.1 per cent reduction in shipping times from 2007 is required to deliver same-day shipping and that 38.2 percentage points of that reduction occurred from 2007 to 2012.

Of course, this massive geographic expansion of warehouses was not without cost. To measure the corresponding increase in cost, I use data on fulfillment expenses as a fraction of the cost of goods sold from Amazon.com's annual reports in 2007 and 2012. Fulfillment expenses incorporate both direct shipping costs paid to courier services as well as amortized investments required for fulfillment and hence are a good measure of the total cost of sending a good to a final customer. From 2007 to 2012, fulfillment expenses increased from 11.3 to 13.9 per cent of Cost of Goods Sold, roughly 5.1 billion dollars (unadjusted for inflation). When calculating welfare gains from reductions in shipping, I include the corresponding 2.6 percentage point increase in fulfillment expenses per purchase.

[14] As of April 2014, same-day delivery is currently available in Baltimore, Boston, Chicago, Indianapolis, Las Vegas, New York City (and parts of New Jersey), Philadelphia, Phoenix, San Bernardino Area, Seattle, and Washington, D.C.

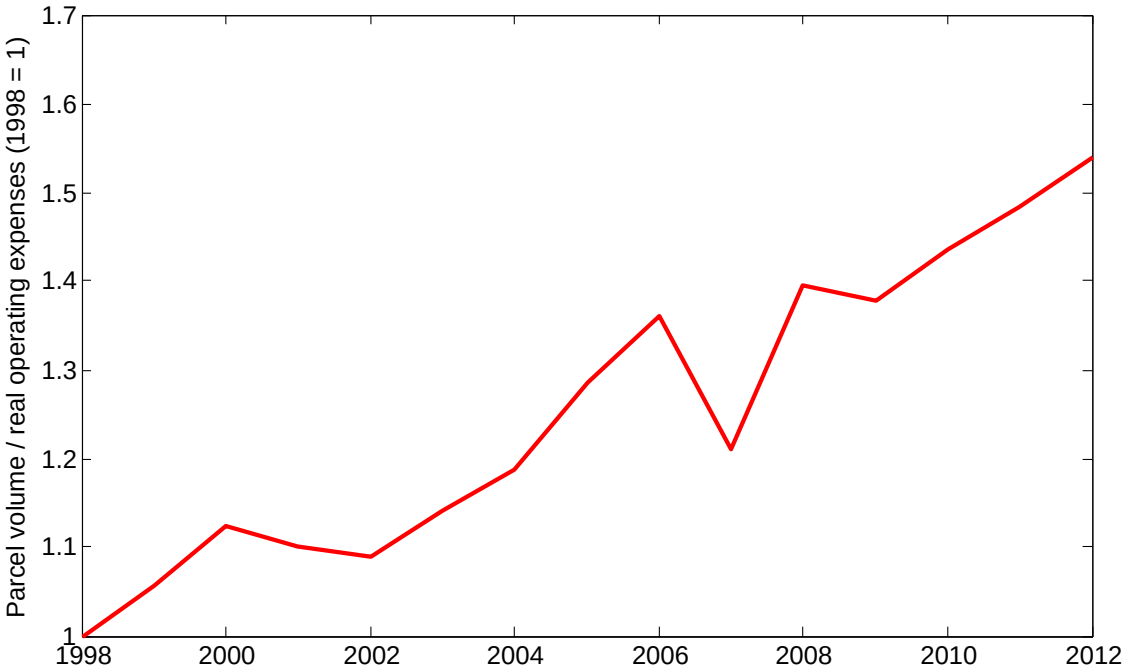

Figure 4: Productivity improvements at FedEx and UPS

5.1.3 Expansion into new industries

Figure 5 shows that the distribution of the online share of sales is concentrated in a few industries. Less than 10 per cent of the 6-digit NAICS industries have an online share of sales greater than 10 per cent. Moreover, very large industries such as supermarkets and grocery stores have virtually non existent online sales. These facts point to potentially large gains for online retailers of expanding into these untapped markets. As anecdotal evidence, consider Amazon.com's entry into the grocery market with Amazon Fresh. Amazon.com seems to exploiting its large shipping and fulfillment network to enter the grocery business, which offers a potentially large source of growth.

As evidence of expansion into new industries, I plot mean annual online sales growth by products between 2007-11 and the log of online sales of products in 2007. Although the relationship is not perfect, there is a negative relationship suggesting that the source of growth in online sales post 2007 has been in the industries with relatively few online sales. The relationship is even stronger if one excludes food, beer and wine which will likely become one of the largest sources of growth for online retailers.

To capture this in the model, I improve online retailers' relative utility, where relative utility is defined as online retailers utility relative to the unweighted mean of utility at brick and mortar stores.

$$\bar{u}_o - \bar{u}_{bm} = \mathbb{E}\left(u_{l,o}\right) - \mathbb{E}\left(u_{l,i}\right) \tag{20}$$

For industries with relatively few online sales, relative utility is small as each visit online by a household generates relatively little surplus compared with a visit at a brick and mortar store. Without a straightforward way to quantitatively measure these improvements, I treat

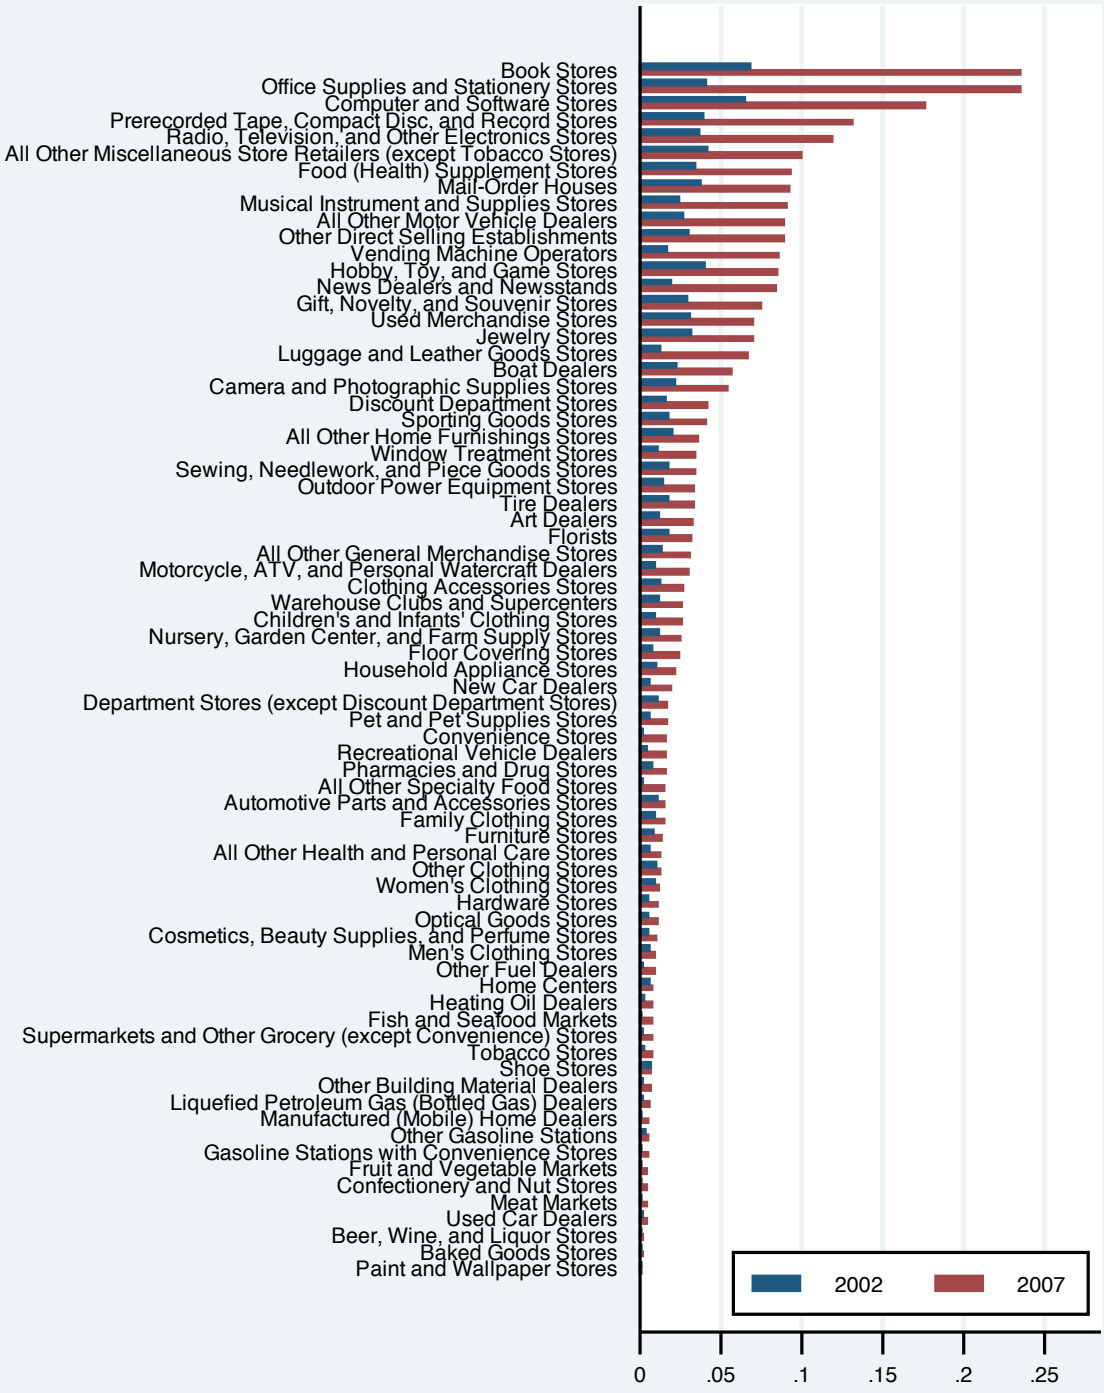

Figure 5: Online market shares by industry

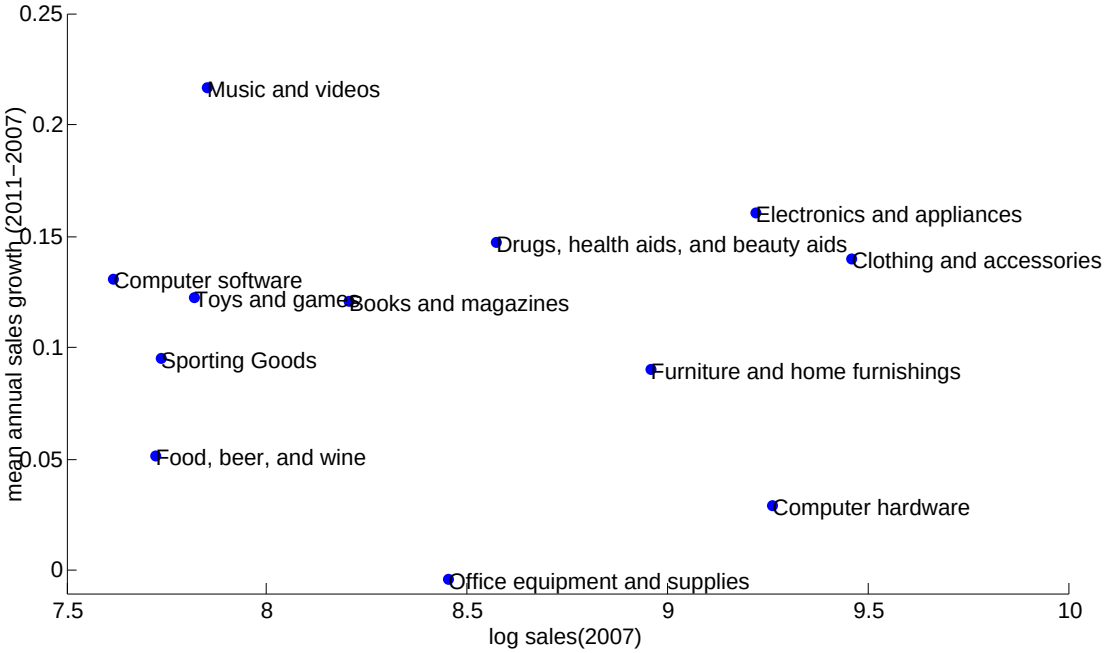

Figure 6: E-commerce sales growth across products (2007 to 2011)

online retailers' expansion into new industries as a residual that soaks up the part of the rise in the online share of sales unaccounted for by increased internet access and reduced shipping times. In particular, I increase the relative utility in industries with below average relative utilities to the value that is necessary for the model to explain the total rise in the online share of sales.

5.1.4 Effects on aggregates

Table 9 summarizes the results from these counterfactuals. I calculate equilibrium outcomes for scenarios where store closure is and is not allowed, and where the size of the shocks mirror those observed in the data and the maximum possible change. Regardless of the counterfactual, effects are largest when one allows stores closures. The rise in the online share of sales increases by between 0.3-0.5 percentage points with store closures. Regardless of exit, improvements in online retail lead some consumers to substitute away from stores to online retailers. However, there is a second bout of reallocation with store closures as the consumers of stores that subsequently close, who did not initially shop online, are forced to find an alternative. From this point on, I focus on the scenario with exit unless stated otherwise.

The results indicate that reduced delivery times and online retailers' expansion into new industries are more important than increased internet access in driving the rise in online sales. More widespread internet access increased the online share of sales by 1.4 percentage points, just over half the impact from reduced shipping times and less than one third relative to

the general expansion into new industries. Combined, increased internet access and reduced shipping times are capable of explaining just under half of the observed increase in the online share of sales.

Going forward in time, a similar ordering applies with the maximum feasible effects. Moving to complete internet penetration raises online sales by 32.9 per cent from 2007 which represents an increase of 14.9 percent from the value estimated for 2012. Shipping improvements are much more effective, increasing online sales by 30 percent from 2012 on. The larger impact of shipping mirrors the online retail industry's focus on ever-faster shipping.

With all three shocks present, the online share of sales increases by 5.2 percentage points. Bear in mind that this exactly matches the observed increase in the online share of sales as the expansion into new industries counterfactual is calibrated to match the observed increase. The aggregate impact on the retail sector is large. Sales at brick and mortar stores fall 15.5 percent with 78,000 stores closing.

$$OX(i) = \sum_{p \in P} \frac{sales_{p,2002}(i)}{sales_{2002}(i)} \left(\frac{sales_{p,2007}(online)}{sales_{p,2002}(online)} - 1 \right) \tag{21}$$

In order to validate these results, I compare the predicted outcomes from the model to data. However, I do not compare them directly to observables since the great recession occurred during the same period. Instead, I compare them to the predicted values from a reduced form model of the relationship between brick and mortar sales and the rise in the online share of sale.

The obvious strategy would be to use variation in online shares of sales and retail sales across industries. The issue with this strategy is that industry based models of the effect of the internet on retail store sales suffer from the problem of confounding effects from the internet with unobservable industry level effects. For instance, sales at book stores may appear to decline because of a corresponding rise in the online share of sales but instead are driven by changing consumer preferences that are independent of online retail. To overcome this problem, I use within-industry variation in the product mix comprising each store's sales.

As an example, consider book stores. The bookstores that support their sales with a relatively large amount of food and beverage sales are relatively insulated from the rise in online sales of books, relative to those that sell only books since food and beverage sales are not subject to online competition whereas books are. The caveat for this identification strategy to be valid is that industry level shocks are not correlated with the product mix of stores. If it were the case that book stores who sell disproportionately more books bore the brunt of changing consumer preferences, then the identification strategy would not be valid.

To measure each store's exposure to online competition, I create an establishment specific online exposure index which is the weighted average of each product's exposure to online retail where the weights come from the composition of the establishment's sales. A product's exposure to online retail is defined as the change in online sales for that product relative to the product's total sales in the initial period. For example, a bookstore with 20 per cent of sales from coffee and 80 per cent books has an online exposure index weighted 20 and 80 percent towards the growth of online sales relative to initial sales of coffee and books respectively. More precisely, let subscript p denote a product and i denote an establishment. The online exposure index is defined as follows.

Table 9: Results from counterfactual exercises

	Observed		**Maximum**	
	No exit	Exit	No exit	Exit
Increased internet access				
Δ online share of sales (ppt)	0.5	0.8	1.1	1.6
Δ online sales (%)	10.8	15.6	23.5	32.9
Δ brick and mortar sales (%)	-1.3	-3.7	-2.8	-6.5
Δ brick and mortar stores (%)	-	-5.1	-	-8.6
Δ brick and mortar stores ('000s)	-	-26	-	-44
Reduced shipping times				
Δ online share of sales (ppt)	1.1	1.4	2.6	3.1
Δ online sales (%)	25.2	29.9	60.0	68.9
Δ brick and mortar sales (%)	-1.3	-3.3	-2.7	-6.1
Δ brick and mortar stores (%)	-	-7.9	-	-12.7
Δ brick and mortar stores ('000s)	-	-40	-	-64
Expansion into new industries[1]				
Δ online share of sales (ppt)	2.4	2.9		
Δ online sales (%)	46.7	57.2		
Δ brick and mortar sales (%)	-6.2	-10.2		
Δ brick and mortar stores (%)	-	-5.3		
Δ brick and mortar stores ('000s)	-	-26		
Expansion into new industries				
+ Reduced shipping times				
Δ online share of sales (ppt)	1.8	2.4	4.3	5.4
Δ online sales (%)	39.1	47.8	96.4	115.1
Δ brick and mortar sales (%)	-2.7	-6.3	-6.0	-7.8
Δ brick and mortar stores (%)	-	-10.9	-	-18.7
Δ brick and mortar stores ('000s)	-	-55-	-95	
Increased internet access				
+ Reduced shipping times				
+ Expansion into new industries				
Δ online share of sales (ppt)	5.2	5.2	8.1	10.7
Δ online sales (%)	95.9	162.2	166.0	208.9
Δ brick and mortar sales (%)	-9.6	-15.5	-13.8	-21.4
Δ brick and mortar stores (%)	-	-15.4	-	-23.2
Δ brick and mortar stores ('000s)	-	-78	-	-118

[1] Expansion into new industries increase online retailers' relative-utility-to-stores in industries with below mean online shares of sales to the mean relative-utility-to-stores.

$$OX(i) = \sum_{p \in P} \frac{sales_{p,2002}(i)}{sales_{2002}(i)} \left(\frac{sales_{p,2007}(online)}{sales_{p,2002}(online)} - 1 \right) \tag{22}$$

To determine the relationship between online retail and brick and mortar sales, I estimate a linear regression model of the log change in a store's sales on the store's online exposure index that takes into account 6 digit industry effects, firm size, population density at the store

location, establishment size and payroll share of sales. Note that since I include industry dummies, any industry level effects are accounted for. The results are provided in Table 10 and as expected, indicate that stores with greater exposure to online competition experienced less sales growth.

Table 10: Regression on store sales

Variable[1]	**Log sales**
constant	0.414 (0.018)
age	-0.006 (0.000)
log(firm size)	0.018 (0.001)
log(pop. density)	-0.012 (0.001)
log(firm payroll / firm sales)	0.110 (0.005)
$OX(i)$	-1.848 (0.343)
$OX(i)\times$log(pop. density)	0.098 (0.027)
$OX(i)\times$ log(firm size)	-0.111 (0.015)

[1] N = 266,000 establishments; Industry dummy variables at the 6 digit level included.

It is straightforward to use the reduced form model to generate a prediction for the reduction in total retail store sales due to online retail. Between 2007 to 2011, average firm size in retail was on 26 employees with average population density in the US 34.2 people per square kilometer. The online exposure index for online retail as a whole in 2007 is 0.064. With these figures, the model predicts that online retail would be responsible for 10.4 per cent decline in retail store sales between 2007 and 2012. In comparison, the structural model's prediction is for a reduction of between 10.5 to 20.5 per cent (lower and upper bounds) in brick and mortar sales which suggests that the structural and reduced form models offer similar predictions.

5.1.5 Effects on Welfare

Although the reduced form and structural model predict similar outcomes, the structural model has the advantage that it is informative about the heterogeneity in effects from online competition and the associated effects on welfare. Given the expression for welfare, it is straightforward to decompose the change in aggregate welfare into direct effects and indirect effects. Letting $\frac{u_{l,i}}{\eta} + p_i - c \equiv z_{l,i}$, the change in aggregate welfare is comprised of: direct changes in welfare, indirect effects from net reallocation and changes in store operating costs.

$$W' - W = \sum_l \sum_i (\Delta y_{l,i} z_{l,i} + y_{l,i} \Delta z_{l,i}) + \sum_i \Delta \varphi_i \tag{23}$$

Net reallocation measures the effect of reallocating households from one store to another and occurs in the model for two reasons. The first is the initial switching by some households to online retailers as utility per visit improves at online retailers. This initial bout of reallocation reduces sales at some stores, leading to store closures and a second bout of reallocation as households who previously shopped at closed stores have to find alternatives. Net reallocation is not necessarily a positive force on welfare as store closures force some households to switch to less preferred alternatives. A more subtle reason why net reallocation may reduce

aggregate welfare is due to measurement and the random component in utility. When online retailers improve, some households shopping at stores offering relatively large consumer surplus receive a large idiosyncratic shock and decide to shop online. Because my measure of aggregate welfare does not capture he random component of utility, this may appear as a household switching to a lower surplus alternative.[15]

Reallocation also leads to store closures and reduces the total cost of operating stores, partially offsetting the negative effect on consumer surplus of store closures.[16] However, note that the reallocation towards online retailers is not without cost. Increased marginal costs from the online retailer are taken from financial statements at Amazon.com and accounted for in the calculation of welfare gains. Marginal costs as measured by operating costs over cost of goods sold are slightly higher at Amazon.com relative to the average retail store.

The final component of aggregate welfare gains come from the direct increase in welfare from improvements by online retailers, holding fixed the allocation of consumers. The total surplus per purchase increases at online retailers as consumer surplus increases more than costs do, ignoring both product and shipping prices.

Table 11: Welfare decomposition

	Observed			
	No exit		Exit	
Increased internet access				
+ Reduced shipping time				
direct welfare gains		5.4		5.7
net reallocation				
non exiters	4.3		3.2	
exiters	-	4.3	-13.1	-9.9
store operating costs		-		12.1
Δ welfare (%)		9.7		8.0
Increased internet access				
+ Reduced shipping time				
+ Expansion into new industries[a]				
direct welfare gains		10.3		8.2
net reallocation				
non exiters	9.2		13.5	
exiters	-	9.2	-21.9	-8.4
store operating costs		-		13.5
Δ welfare (%)		19.5		13.4

[a] Expansion into new industries such that all counterfactuals with all three improvements matches observed increase in online share.

Table 11 shows the effects on aggregate welfare from these counterfactuals. Consider the results from the counterfactual with both the observed increase in internet access and reductions in shipping time. These observed changes can account for roughly half of the

[15] This is analogous to entry by a McDonalds store offering terrible service stealing customers from all restaurants in a location, rather than just the poorly performing ones.

[16] Of course, this represents a transfer of surplus if ownership of firms is unequal across households.

increase in the online share of sales and lead to welfare gains of 8 per cent. Results including online retailers' expansion into new industries are much larger with welfare estimated to increase by 13.4 per cent, with the bulk of these gains coming the direct component of welfare gains.

These results highlight the substantial dampening effect store closures can have on aggregate welfare. Exiters subtract between 13.1 to 21.9 percentage points off welfare gains with the subsequent reallocation to second best alternatives recuperating only 3.2 and 13.5 percentage points. The substantial negative effect of net reallocation indicates that there are many consumers who actually lose out from improvements by online retailers. Overall, these store closures and reallocation leads to positive effects on welfare as there are savings in the costs of operating stores. However, these savings are captured by owners of firms which are unlikely to compensate consumers unless they are owners of retail firms.

5.1.6 Effects on firms

Table 12 decomposes the aggregate change in retail profits caused by improvements in online retail. The counterfactuals suggest that aggregate profits decline by 8.2 per cent. The aggregate decline masks divergent outcomes for brick and mortar firms compared to online retailers. Online retailers add 8.8 percentage points to aggregate profits which is considerable given their size while brick and mortar profits subtract 17 percentage points. The fall in brick and mortar profits is driven by losses at both exiters and non exiters, who subtract 3 and 14 percentage points off aggregate growth respectively.

Table 12: Decomposing aggregate profits

	Observed			
	No exit		Exit	
Increased internet access				
+ Reduced shipping time				
+ Expansion into new industries				
online retailers		7.8		8.8
brick and mortar stores				
non exiters	-23.2		-14.0	
exiters	-	-23.2	-3.0	-17.0
Δ total profits (%)		-15.4		-8.2

Although improvements by online retailers unequivocally lead to less sales at traditional retailers and store closures, effects on retail firms are heterogeneous since competitive pressure varies across areas. Stores that manage to survive can often experience an increase in profits if they soak up enough market share from exiting stores to offset the loss of customers to online retailers. When increased competition from online retail induces a store to close, former customers of the closed store are reallocated to a variety of stores, not only to online retailers. Stores that have relatively fat operating margins are less likely to exit from an initial bout of increased competition, and stand to gain if there are a large number of store closures nearby.

To gauge the extent to which some firms are better off, I calculate percentiles of the percentage change in firm profits across surviving firms. The results are displayed in Table 13. The median surviving firm suffers a 12.2 reduction in profits with the top 10 percent

Table 13: Distribution of profit changes at surviving firms

	Percentiles - change in profits (%)						
Average across industries	**5**	**10**	**25**	**50**	**75**	**90**	**95**
Increased internet access							
+ Reduced shipping time							
+ Expansion into new industries							
No exit	-100.0	-100.0	-100.0	-54.2	-17.1	-5.2	-2.6
Exit	-69.6	-54.3	-30.2	-12.2	-3.0	0.1	20.2

of surviving firms actually experiencing an increase in profits. Because 17.9 percent of firms exit with all three shocks, the model suggests that 8.2 percent of firms increase profits in response to improvements by online retails. Moreover, these figures are likely to understate the increase in profits at some firms as the model does not account for increased online sales by previously brick and mortar-only firms.

5.2 Effects of the Marketplace Fairness Act (2013)

Because of a 1967 Supreme Court case, out-of-state retailers have been exempt from collecting state sales taxes from customers due to the complexity of collecting taxes for multiple states, with the many variations in tax rates, exemptions and record keeping requirements. The Marketplace Fairness Act, before Congress at the time of publication, seeks to overcome this ruling by allowing states to compel online retailers to collect state taxes in compensation for the simplification of states sales tax laws.

Since state sales taxes are already collected by traditional retailers, the passing of the act would represent an increase in the relative price of goods purchased online. Supporters therefore claim that the act would level the playing field between online retailers and traditional retailers. To simulate the implementation of the act, I simply increase the effective markup that online retailers charge over marginal costs by the sales tax rate based on the purchasing customer's location. States are required to establish a uniform tax rate before they can compel online retailers to collect sales taxes. For this reason, I use sales tax rates based on the state tax rate and the average combined city and county tax rates as of August 2013. A summary of these tax rates are listed in Table 14.

Table 14: Average state tax rates inclusive of city and county rates

Tax rate (%)	**States**
≤ 5	Alaska, Delaware, Hawaii, Maine, Montana, New Hampshire, Oregon
>5 & ≤ 6	District of Columbia, Kentucky, Maryland, Michigan, Nebraska, North Dakota, South Dakota, Virginia, Wisconsin, Wyoming
>6 & ≤ 7	Colorado, Connecticut, Florida, Georgia, Idaho, Iowa, Indiana, Massachusetts, New Jersey, New Mexico, North Carolina, Pennsylvania, Rhode Island, Utah, Vermont, West Virginia
>7 & ≤ 8	Arizona, Ohio, Kansas, South Carolina, Texas
> 8	Arkansas, California, Oklahoma, Illinois, Louisiana, Tennessee, Washington

The estimates for the utility from the online option do not include the contribution of

markups since it was not necessary to estimate the components separately. To get the price of the online good from which to implement the sales taxes, I use data from the financial statements of Amazon.com to calculate markups inclusive of shipping to get measures of net markups for the online sector in 2012.[17] Gross margins net of shipping discounts at Amazon.com were 20.1 per cent in 2012 which implies a value for p_o of 1.252. Given this estimate of the price of the online good, the utility from shopping online becomes $u_{o,tax}$ with the tax rate, τ, dependent on the location of the purchasing customer.

$$u_{o,tax} = u_{o,t} - \eta p_o \tau$$

Before imposing the online sales tax, I extrapolate the model to 2012 by using the counterfactual where internet access improves, delivery times fall and there is a general expansion of online retail across newer retail categories. Therefore results should be interpreted as the predicted outcome had the MFA been implemented sometime during 2007-12.

Table 15: Estimated potential effects of Marketplace Fairness Act (2007-12)

	No exit		**Exit**	
	Baseline	**Tax**	**Baseline**	**Tax**
Δ online share of sales (ppt)	5.2	2.6	5.2	3.4
Δ online sales (%)	95.9	58.0	162.2	69.9
Δ brick and mortar sales (%)	-9.6	4.8	-15.5	-9.3
Δ brick and mortar stores (%)	-	-	-15.4	-11.9
Δ brick and mortar stores ('000s)	-	-	-78	-60

The results from are displayed in Table 15. Absent implementation of the tax, online sales increase by between 162.2 per cent with 78,000 store closures. The tax has a large effect, paring increases in online sales by roughly 90 percentage points leading to online sales in 2012 that are 0.54 times online sales without the tax. These numbers are larger than estimated from a price comparison website in Ellison and Ellison (2009), which calculates a decline in online sales of 30 per cent if offline sales taxes were eliminated. It is not surprising that the impact here is larger since the data in Ellison and Ellison (2009) come from a price comparison website for memory modules, where price sensitivity is likely to be relatively high.

6 Conclusion

In this paper, I estimated that online retail had a large positive effect on welfare. The estimated model suggests that the 5.2% increase in the online share of sales from 2007-12 is associated with a 13.4% increase in welfare, which incorporates both producer and consumer surplus. This magnitude of these effects have broader implications because the Retail and Wholesale Trade sectors are large, and also because they are indicative of further gains in other sectors being transformed by online services.[18] The welfare gains from this internet technology

[17] Domestic sales at Amazon.com represented roughly 13 per cent of total online sales in 2011.

[18] For example, consider the rapid growth of Uber, Lyft and Sidecar in transportation and Square in financial services.

are in addition to the prior productivity improvements in logistics and inventory management due to information technology. In all, these technological advances and their effects on welfare are strong evidence of the role of recent advances in computer-based technology in spurring economic growth.

Key to quantifying the effects on welfare is measuring each store's contribution to welfare, akin to store-level retail productivity, accurately. Instead of using standard measures of retail productivity such as value added per worker to measure a store's contribution to welfare, I use a measure that reflects the surplus (consumer surplus and profits) that each visit by a consumer generates. This new measure is robust to variation in competitive pressure, unlike standard measures of productivity which are distorted. Of course, the tradeoff is that a measure of competition is required for every store, which I obtain by using a structural model.

I focus on two margins by which online retail affects aggregate welfare in the retail sector. The first are the gains in consumer surplus from improvements in online retail holding fixed market shares. The next is the change in the market structure of retail induced by reallocation. The changing market structure is complicated as reallocation occurs locally rather than at an aggregate level and the bulk of reallocation is achieved by store closures. Reallocation incorporates the optimal switching of consumers from stores to online retail, the negative effect on consumer surplus from store closures and the savings in store operating costs from closing stores.

To account for these rich local-level selection dynamics and to measure store-level economic surplus, I develop a model of retail that is estimated on store-level data spanning the universe of retail stores in the US. The model takes the geography of the US seriously with stores existing in locations that reflect their real-world location, with competition for customers occurring at a local level. Crucially, the model is one of industry equilibrium with stores setting prices based on local market power and store closures based on firm profit maximization. Therefore, counterfactuals can be used to identify the effects on aggregate welfare, holding all else constant.

Estimates from the model exhibit substantial within-industry heterogeneity in stores' fixed costs and distributions of store quality (economic surplus per store) across locations. The extent of heterogeneity in fixed costs implies that selection is based on store profitability rather than on store quality alone. Moreover, since distributions of store quality vary across locations, selection occurs between stores operating in local markets rather than between retail firms competing in an aggregate market. Selection dynamics are markedly different from selection effects in a typical heterogeneous firms model, where the lowest productivity establishments are the first to exit in response to an adverse shock. Stores offering consumers relatively little rather than high quality stores may survive if they are more profitable. In counterfactual exercises, a large fraction of firms actually experience an increase in profits from improvements by online retailers as they soak up enough market share from exiting stores to offset lost market share to online retailers.

The counterfactual exercises assess the effect on aggregate welfare from 3 potential drivers of the increase in online sales. The observed increase in internet access from 2007 to 2012 is estimated to increase the online share of sales by 0.8 percentage points. Reduced shipping times that reflect Amazon.com's geographic expansion of its shipping and fulfillment warehouses and productivity improvements at UPS and FedEx from 2007 to 2012 are estimated to have a much larger effect, increasing the online share by 1.4 percentage points. The largest

effect comes from expansion into new industries, which increases the online share by 2.9 percentage points. Results from a reduced form model confirm the aggregate decline in brick and mortar sales suggested by the structural model. The cumulative effect of these 3 changes increase welfare by 13.4%. Net reallocation subtracts 8.4 percentage points off the aggregate figure, while direct welfare gains and savings in store operating costs add 8.2 and 13.5 points respectively.

Finally, the paper considers the effect of the Marketplace Fairness Act, a bill before Congress that would allow states to compel online retailers to collect state sales taxes. I introduce that the tax as a state specific distortion to relative prices. The model estimates a large effect from the tax if implemented between 2007-12. Sales at brick and mortar stores would have decreased by one-third less, saving 18,000 stores in the process. The results suggest that online sales would fall by 46 percent, slightly higher than the figure of 30 percent estimated in Ellison and Ellison (2009) which uses data from a price comparison site for memory modules.

References

Abbring, J. H., Campbell, J. R., and Yang, N. (2012). Simple Markov-Perfect Industry Dynamics.

Arkolakis, C. (2010). Market penetration costs and the new consumers margin in international trade. *Journal of Political Economy*, 118(6).

Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile Prices in Market Equilibrium. *Econometrica*, 63(4):841–890.

Bresnahan, T. F. and Reiss, P. C. (1991). Entry and Competition in Concentrated Markets. *Journal of Political Economy*, 99(5):977–1009.

Campbell, J. R. and Hopenhayn, H. A. (2005). Market Size Matters. *The Journal of Industrial Economics*, 53(1):1–25.

Chiou, L. (2009). Empirical Analysis of Competition between Wal-Mart and Other Retail Channels. *Journal of Economics & Management Strategy*, 18(2):285–322.

Collard-Wexler, A. and Loecker, J. D. (2013). Reallocation and Technology: Evidence from the U.S. Steel Industry. Technical Report 18739.

Davis, P. (2006). Spatial Competition in Retail Markets: Movie Theaters. *The RAND Journal of Economics*, 37(4):964–982.

Ellison, G. and Ellison, S. F. (2009). Tax Sensitivity and Home State Preferences in Internet Purchasing. *American Economic Journal: Economic Policy*, 1(2):53–71.

Foster, L., Haltiwanger, J., and Krizan, C. J. (2006). Market Selection, Reallocation, and Restructuring in the U.S. Retail Trade Sector in the 1990s. *Review of Economics and Statistics*, 88(4):748–758.

Holmes, T. J. (2011). The Diffusion of Wal-Mart and Economies of Density. *Econometrica*, 79(1):253–302.

Kaplan, G. and Menzio, G. (2013). Five Facts About Price Dispersion.

Lagakos, D. (2013). Explaining cross-country productivity differences in retail trade.

Morales, E., Sheu, G., and Zahler, A. (2013). Gravity and Extended Gravity: Estimating a Structural Model of Export Entry.

Oi, W. (1992). Productivity in the Distributive Trades: The Shopper and the Economies of Massed Reserves. In *Output Measurement in the Service Sectors*. NBER.

7 Appendix

7.1 Estimates of internet access by zip code

To obtain estimates of internet access, I use microdata from the CPS. Unfortunately, the finest geographic level at which the data are available is at the MSA level which is relatively coarse and covers only metropolitan areas. However, the CPS provides data on each respondent to the survey, with weights that aggregate up to the population level. This demographic data allows me to generate zip code level estimates of internet access by applying estimates from individual level models to more aggregated data. First, I use demographic and geographic data on each respondent to estimate a model of individual level internet access. Given these estimates, I construct estimates for the fraction of the population with access to the internet by zip code, by applying the model to demographic data for each zip code. This is equivalent to assuming that representative individuals populate each zip code, each characterized by zip code-level demographic data.

Table 16: Internet access and demographics

Variables[1]	**1997**	**2002**	**2007**	**2012**
constant	-14.77	-14.15	-11.82	-9.40
age	0.0022	-0.0008	0.0012	0.0029
age^2	-0.0003	-0.0003	-0.0003	-0.0003
log(household income)	1.30	1.33	1.22	1.06
Dummy Variables				
white	0.28	0.74	0.39	0.33
black	-0.66	-0.11	-0.26	-0.21
asian	0.32	1.01	0.49	0.49
hispanic	-0.68	-0.29	-0.60	-0.36
high school	-0.05	0.08	0.14	0.27
some college	0.48	0.58	0.74	0.72
associates degree	0.46	0.57	0.78	0.81
bachelors degree or higher	0.85	0.98	1.15	1.25

[1] N $\equiv$ min. 270 miliion; includes state dummy variables. Race and education groups are mutually exclusive.

The results are broadly consistent with priors that internet access increases with income and education, decreases with age with whites and asians having an advantage over hispanics, blacks and others. As expected, the model suggests that the fraction of the population with internet access has increased over time which is consistent with aggregate CPS data which indicate that the percentage of the population with internet access at home was 29.9, 56.6, 75.6 and 85.8 in 1998, 2003, 2007 and 2012 respectively. Increased access to the internet country wide is reflected in the estimate of the constant parameter which is increasing monotonically with time. While internet access has increased overall, the model also indicates that the increase has been relatively larger for some groups, likely reflecting saturation. In 1997, a white individual had 2.55 and 2.61 times the odds of having access to the internet compared to black and hispanic individuals. By 2012, the differences in the odds of internet access for whites compared to blacks and hispanics had dropped to 1.71 and 1.99 respectively as

relatively more blacks and hispanics obtained access to the internet. Similarly, household income becomes less important over time as a predictor of internet access. Note that this is not driven by inflation since I deflate household income using the CPI to 2007 dollars.

To extrapolate these individual level estimates to the zip code level, I require zip code level data on all the variables that the individual level model was estimated on. All the variables map to variables available by zip code in the 2000 Decennial Census and the 5 year American Community Survey (ACS) 2007-2011. I use data from the 2000 Decennial Census for years 1997 and 2002 and data from the ACS for the years 2007 and 2012. Results have also been generated using linear interpolation/extrapolation but do not generate any meaningful differences and hence are not reported. To obtain estimates of internet access by zip code for a given year, I apply the individual level model with coefficients from the given year to demographic data by zip code for that year. The key difference is that rather that having dummy variables, I have proportions of a zip code's population that fit each education and race group. Hence the estimates are akin to estimating the probability that a representative individual has access to the internet for each location.

7.2 Proof of uniqueness and existence

For clarity, I prove uniqueness and existence for the model with complete internet access. The proof of the model with heterogeneous internet access is a straightforward extension where weights of market shares need to be accounted for. The proof that equilibrium prices exist and are unique relies on proving that the following operator is a contraction. Define the operator $T : \mathbb{R}^k \to \mathbb{R}^k$, element by element as follows

$$T_j(p) = p_j + \beta \left(\log \left(c + \frac{1}{\eta \left(1 - \bar{v}_i(p)\right)} \right) - \log(p_j) \right) \tag{24}$$

for some $\beta \in (0, 1)$ and $\bar{v}_i(p)$ is calculated using p.

$$\bar{v}_i(p) = \sum_{l \in \mathbb{H}(i)} \omega_l v_{l,i}(p)$$

where

$$\omega_l = \frac{m(l)n(l)}{\sum_{l \in \mathbb{H}(i)} m(l)n(l)}$$

Lemma 1. *$\exists\ \beta$ such that $\forall j, k$, $\frac{\partial T_j(p)}{\partial p_k} \geq 0$ and $\sum_k \frac{\partial T_j(p)}{\partial p_k} < 1$.*

Proof. Let i be an arbitrary store and denote $v_{l,max}(i)$ as the market share in location l for store i when $p_i = c$ and $\forall j \neq i, p_j \to \infty$. Define the maximum market share as follows $\bar{v} = \max_i \max_l v_{l,max}(i)$. Note that $\bar{v} < 1$ due to the presence of the outside options, which always have positive market share due to the logit error. Taking derivatives of equation (24),

$$\frac{\partial T_j(p)}{\partial p_j} = 1 - \beta \left(\frac{\sum_l \frac{\omega_l (1 - v_{j,l})}{1 - \bar{v}_j} v_{j,l}}{c(1 - \bar{v}_j) + \frac{1}{\eta}} + \frac{1}{p_j} \right)$$

and $\forall k \neq j$,

$$\sum_k \frac{\partial T_j(p)}{\partial p_k} = \beta \frac{\sum_l \sum_k \frac{\omega_l v_{j,l}}{1-\bar{v}_j} v_{k,l}}{c(1-\bar{v}_j) + \frac{1}{\eta}}$$

It suffices to show that $\frac{\partial T_j(p)}{\partial p_j} \geq 0$ and that $\sum_k \frac{\partial T_j(p)}{\partial p_k} < 1$ which in turn reduce to showing that the following hold.

$$\beta \left(\frac{\sum_l \frac{\omega_l (1-v_{j,l})}{1-\bar{v}_j} v_{j,l}}{c(1-\bar{v}_j) + \frac{1}{\eta}} + \frac{1}{p_j} \right) \leq 1 \tag{25}$$

$$\frac{\sum_l \sum_k \frac{\omega_l v_{j,l}}{1-\bar{v}_j} v_{k,l}}{c(1-\bar{v}_j) + \frac{1}{\eta}} < \frac{\sum_l \frac{\omega_l (1-v_{j,l})}{1-\bar{v}_j} v_{j,l}}{c(1-\bar{v}_j) + \frac{1}{\eta}} + \frac{1}{p_j} \tag{26}$$

To ensure the equation (25) holds, one can simply set β low enough such that equation (25) binds for a value for the term inside the brackets guaranteed to be greater than for any store. Setting $\beta = \beta_{bound}$ ensures equation (25) is satisfied for all stores as $\forall j, p_j > c, v_{j,l} < 1$ and $\bar{v}_j < \bar{v}$.

$$\frac{1}{\beta_{bound}} = \frac{\sum_l \frac{\omega_l}{1-\bar{v}}}{c(1-\bar{v}) + \frac{1}{\eta}} + \frac{1}{c}$$

Finally, equation (26) holds since prices are positive, $p_j > 0$, and the presence of the outside options guarantee that $\sum_k v_{k,l} < 1 - v_{j,l}$.

□

Lemma 2. *There is a value, $\bar{p}$, such that if for any j, $p_j > \bar{p}$, then for some k, $T_k(p) < \bar{p}$.*

Proof. Set $\bar{p} = c + \frac{1}{\eta(1-\bar{v})} + \epsilon$ for any $\epsilon > 0$ where $\bar{v}$ is defined in the previous lemma. Let j be an arbitrary store and assume $p_j > \bar{p}$. It suffices to show that $\log\left(c + \frac{1}{\eta(1-\bar{v}(i;p))}\right) - \log(p_j) < 0$.

$$\begin{aligned} \log\left(c + \frac{1}{\eta\left(1-\bar{v}(i;p)\right)}\right) &\leq \log\left(c + \frac{1}{\eta\left(1-\bar{v}\right)}\right) \\ &< \log\left(p_j\right) \end{aligned}$$

□

Proof of existence and uniqueness

Proof. The operator T satisfies the conditions of the contraction proof in Berry et al. (1995) as Lemmas 1 and 2 hold. Since T is a contraction, there exists a unique fixed point to the operator T. This implies that the equilibrium defined by (9) exists and is unique since the existence of another equilibrium would contradict the uniqueness of the fixed point to the operator T. □

7.3 BLP algorithm with endogenous prices

Guess initial values for each store, $\left\{\hat{z}_j^0\right\}$. For every store and all its customers, calculate the probability that a customer in a location chooses that store assuming that for all brick and

mortar stores prices are zero and the $\left\{\hat{z}_j^0\right\}$ are the unobserved store effects.

$$\hat{v}_l(i) = \frac{e^{\theta_{f(i)} + \mathbf{x_i}'\theta_i - \left(\zeta_0 + \zeta_1 \log popden_{i,l}\right) d_l(i) + \hat{z}_i}}{1 + e^{u_{o,t}} + \sum_{j \in \mathbb{S}(l)} e^{u_l(j)}} \tag{27}$$

Calculate prices according to the following equation

$$p(i) = c + \frac{1}{\eta\left(1 - \hat{v}(i)\right)} \tag{28}$$

where $\hat{v}(i)$ represents a store's sales-weighted average market share using equation (27).

$$\hat{v}(i) = \sum_{l \in \mathbb{H}(i)} \frac{y_l(i)}{\sum_{l' \in \mathbb{H}(i)} y_{l'}(i)} v_l(i)$$

Use equations (27) and (28) to calculate a store's nominal sales to customers in a given location and aggregate over locations to get total sales for each store. For each location, calculate total sales for stores in that location and each store's share of the total. Update the initial guess of $\left\{\hat{z}_j^0\right\}$ with the following updating equation.

$$z_i^k = z_i^{k-1} + \log(s_{i,data}) - \log(s_{i,model}) \tag{29}$$

Repeat the above until the algorithm converges to a fixed point.

Lemma 3. *A store's share of nominal sales calculated according to the algorithm is decreasing in its own $\hat{z}_i$.*

Proof. By using equation (27), it is straightforward to show that the probability of a customer visiting the store falls relative to all other stores in the same location, implying that the store's share of customers fall. Equation (28) shows that the store's price relative to other stores' prices in the same location falls as the store's relative share of customers fall. Finally, stores in the same location compete across a common set of markets implying that the store's relative prices and quantities fall in all markets. □

Theorem 2. *The operator defined by the BLP algorithm with endogenous prices is (i) a contraction, (ii), generates prices that are optimal and (iii) finds store level fixed effects that match observed market shares.*

Proof. Because Lemma 3 holds, the standard proof that the BLP algorithm is a contraction is applicable (Berry et al., 1995). To show that the algorithm generates optimal prices, calculate the store-level fixed effect by transforming the solution to the fixed point problem with the algorithm-generated prices for each store.

$$z_j = \hat{z}_j + \eta p_j \tag{30}$$

By definition, the fixed point $\{\hat{z}_j\}$ and associated prices $\{p_j\}$ satisfies equations (27) and (28). Use (30) to substitute out $\{\hat{z}_j\}$ in equations (27) and (28). Define $\{z_j\}$ as the store level fixed effect. With this definition, it is clear that prices are optimal as the modified

pricing equation, (28), resembles the model pricing equation, (9). Finally, $\{z_j\}$ and $\{p_j\}$ generate market shares that match observed market shares since $\{\hat{z}_j\}$ is a fixed point of the operator. □

7.4 Entry and exit algorithm

It is likely that multiple equilibria exist given that variations in the order of exit lead to different equilibria.[19] As an example, suppose that Store A and Store B (owned by separate firms) exist in some location with no other competitors and that a sudden increase in internet access amongst customers in that location reallocates sales away from Stores A and B. If it is the case that both would want to close if the other remained open and both would want to remain open if the other closed, then two equilibria exist, characterized by either store remaining open. I focus on the equilibrium generated by the following iterative procedure.

Algorithm:

1. Exit step.
   (a) For each store in operation, calculate the value to the firm of closing that store assuming that all other open stores remain open while accounting for equilibrium price changes.
   (b) Rank those stores wishing to close in order of the absolute size of losses.
   (c) Close the store making the largest loss. If no stores wish to close, stop.
   (d) Return to step 1a.
2. Entry step.
   (a) For each store in operation, calculate the value to the firm of opening another store in the same location holding fixed the set of other stores while accounting for equilibrium price changes.
   (b) Rank those stores wishing to expand/re-open in order of the absolute size of gross profits.
   (c) Open the store with the largest potential profit. If no stores wish to open, stop.
   (d) Return to step 2a.
3. If both exit step and entry step conclude without changes in set of stores, stop otherwise return to step 1.

The algorithm has a natural interpretation where the industry alternates between an exit stage and an entry stage and concludes when no store wishes to exit or enter.. In the exit stage, stores exit sequentially in order of those making the largest losses. Stores exit sequentially with the ordering recalculated with each exit. When the exit step concludes, no store in operation wishes to exit. However, this situation is not necessarily an equilibrium as exit of some stores may make entry by other firms optimal and moreover, subsequent exits in the exit step may make previous exits suboptimal. Hence the algorithm needs to account for new store openings and re-entry by stores that exited prematurely in the exit step. The entry

[19] Note that this problem does not arise in estimation since I assume that fixed costs are such that the equilibrium generates the observed set of stores. When performing counterfactuals, there is no such observed set of stores.

step ranks potential firm profits from opening a new store/re-opening a previously closed store and opens stores sequentially until no more firms wish to enter/re-enter. By definition, if both the entry and exit steps conclude without changes in the set of stores, the current set of stores is an equilibrium since no firm wishes to adjust on the extensive margin.

As an example, it is illustrative to see how this resolves the dilemma above. Calculate the losess for each store assuming both stores remain open. Close the store with the largest loss, say, Store B. By definition, the remaining store, Store A, is now profitable. In the entry step, the firm owning Store B by definition does not wish to re-open. For simplicity, assume that the firm which owns Store A does not wish to open a new store. Redoing the exit and entry steps leads to no changes and hence we are left with the equilibrium where Store A remains open, the natural equilibrium where the store making the largest losses closes first or the store making the largest profits remains open.[20]

7.5 Results at bounds of fixed costs

[20]The equilibrium is similar in spirit to the "natural equilibrium" in Abbring et al. (2012).

Table 17: Results from counterfactual exercises

	Observed		**Maximum**	
	Low FC	High FC	Low FC	High FC
Increased internet access				
Δ online share of sales (ppt)	0.5	1.2	1.1	2.1
Δ online sales (%)	11.2	20.0	25.1	40.8
Δ brick and mortar sales (%)	-1.3	-6.1	-3.1	-10.0
Δ brick and mortar stores (%)	-4.0	-6.1	-7.4	-9.9
Δ brick and mortar stores ('000s)	-20	-31	-38	-50
Reduced shipping times				
Δ online share of sales (ppt)	1.2	1.7	2.7	3.5
Δ online sales (%)	26.5	33.4	62.5	74.3
Δ brick and mortar sales (%)	-1.4	- 5.3	-3.1	-9.1
Δ brick and mortar stores (%)	-6.8	-9.1	-12.4	-13.0
Δ brick and mortar stores ('000s)	-35	-46	-62	- 67
Expansion into new industries[1]				
Δ online share of sales (ppt)	2.5	3.8		
Δ online sales (%)	49.3	65.1		
Δ brick and mortar sales (%)	-6.5	-14.0		
Δ brick and mortar stores (%)	-4.9	-5.6		
Δ brick and mortar stores ('000s)	-25	-28		
Expansion into new industries				
+ Reduced shipping times				
Δ online share of sales (ppt)	1.9	3.0	4.5	6.4
Δ online sales (%)	40.8	54.9	102.4	127.8
Δ brick and mortar sales (%)	-2.9	-9.8	-6.9	-16.4
Δ brick and mortar stores (%)	-9.9	-12.0	-18.2	-19.2
Δ brick and mortar stores ('000s)	-50	-61	-93	-97
Increased internet access				
+ Reduced shipping times				
+ Expansion into new industries				
Δ online share of sales (ppt)	5.2	5.2	8.9	12.6
Δ online sales (%)	102.5	129.4	180.2	219.7
Δ brick and mortar sales (%)	-10.5	-20.5	-16.0	-26.9
Δ brick and mortar stores (%)	-14.6	-16.2	-23.2	-23.3
Δ brick and mortar stores ('000s)	-74	-82	-118	-118

[1] Expansion into new industries increase online retailers' relative-utility-to-stores in industries with below mean online shares of sales to the mean relative-utility-to-stores.