Economies of Density in E-Commerce: A Study of Amazon's Fulfillment Center Network *

Jean-François Houde
Department of Economics
Cornell University

Peter Newberry
Department of Economics
The Pennsylvania State University

Katja Seim
The Wharton School
University of Pennsylvania

December 26, 2016

Abstract

We examine the economies of density associated with the expansion of Amazon's distribution network form 2006-2018. We first demonstrate that, in placing a fulfillment center in a new state, Amazon faces a trade-off between the revenue considerations from exposing local customers to sales tax and the cost savings from reducing the shipping distance to those customers. Using detailed data on online transactions, we estimate a model of demand for retail goods and show that consumers' online shopping is sensitive to being charged sales tax. We then use the demand estimates and the spatial distribution of demand relative to Amazon's fulfillment centers to produce predicted revenues and shipping distances under the observed fulfillment center roll-out and under counterfactual roll-outs over this time period. Using a moment inequalities approach, we infer the cost savings associated with being closer to customers that render the observed network roll-out optimal. We find that Amazon saves between \$0.58 and \$1.55 for every 100 miles of \$100 of goods shipped. Further, we calculate the cost saving associated with the expansion of the network over the last decade and find that Amazon has reduced its total shipping cost by over 50% and increased its profit margin by between 5 and 14%. We also demonstrate that prices on Amazon fall by approximately 40%, which mirrors the more than 50% decrease in the average cost of shipping an order.

Keywords: e-commerce, sales tax, Amazon, distribution, logistics, shipping
JEL Codes: H71, L11, L81

*All correspondence may be addressed to the authors via e-mail at houde@cornell.edu, pwnewberry@psu.edu, or kseim@wharton.upenn.edu.

1 Introduction

Online retail has grown substantially over the last decade and now makes up nearly 7% of all retail sales.[1] Amazon.com is a key contributor of this growth, with net sales increasing from $8.5 billion in 2005 to $89 billion in 2014.[2] Amazon's rise has been accompanied by a substantial increase in concentration in online retail, as Figure 1 illustrates. This trend counters initial expectations that the internet would facilitate highly competitive markets due to lower search costs, lower fixed costs on the part of sellers for whom the geography of their physical locations is irrelevant to serving customers, and an increasing reliance on an already existing network of distribution and shipping companies whose services could at least partially replace the offerings of brick-and-mortar stores. In this paper, we explore the possible sources for this increase in concentration with a focus on the competitive advantage that comes with the cost savings of a large-scale distribution network.

We investigate the role of economies of scale in distribution using Amazon as a case study. As Amazon's revenue has grown, so has its network of distribution centers (called Fulfillment Centers, denoted as FCs going forward). The number of FCs has grown from 8 centrally located centers in 2006 to 54 FCs spread out across the US by the end of 2015.[3] Amazon has plans to further expand its network to over 100 facilities by 2018, including "Amazon Now" hubs that provide same-day delivery to local customers. A larger distribution network has implications for firm profitability on both the supply and demand side.

On the supply side, an additional fulfillment center adds fixed costs of opening and operating the facility, relative to a smaller network. Adding a distribution center also makes the existing distribution network more dispersed, shortening the distance to a subset of customers. This can translate into shipping cost savings in one of two ways. First, Amazon's distribution facilities may be closer, on average, to its shippers' local sorting facilities, reducing the nodes through which packages travel in the shipment process. This reduces the shippers' cost, providing Amazon with bargaining power in negotiations with shippers over the per-package rates they charge. Second, by being closer to its customers, Amazon has increasingly been able to enter into the local delivery market itself, replacing traditional shippers in last-leg deliveries in large urban markets near its distribution centers. The extent to which these economies of density affect Amazon's overall profitability is likely significant as shipping comprises a large share of Amazon's operating costs.[4, 5]

The expansion of the network also has the potential to affect consumer willingness-to-pay

[1] www.census.gov/retail/mrts/www/data/pdf/ec_current.pdf.

[2] www.statista.com/statistics/266282/annual-net-revenue-of-amazoncom/.

[3] These figures do not include Amazon fresh FCs or sortation centers. See Section 2 for details on distribution center types.

[4] The financial statements report the costs (net of shipping revenue) are between 3-5% of net sales. These costs totaled nearly $4.3 billion in 2014. See annual reports at http://phx.corporate-ir.net/phoenix.zhtml?c=97664&p=irol-reportsannual.

[5] Expansion also affects the distance of "inbound" shipments from the supplier to Amazon's FC. As we discuss below, suppliers are shipping goods to brick and mortar retail outlets all around the country; as a result, we assume that these costs do not change significantly as Amazon expands its network.

through two opposing demand-side forces. First, if consumers value the convenience of receiving packages faster and expansion leads to shorter delivery times, then such convenience effects constitute a local vertical element to Amazon's service that may in part be responsible for its increasing market share. Second, expansion may increase tax inclusive prices on Amazon through a federal tax law that requires an online retailer to collect sales tax on the customer's behalf *only* if the firm has a physical presence (or "nexus"), such as a warehouse, in the customer's state of residence.[6] As earlier literature (Einav et al. (2014), Baugh et al. (2014)) shows, demand is sensitive to sales taxes as they raise the effective price of an online transaction. To limit the revenue implications of exposing customers to new sales tax liabilities, a supply-side response would thus be not to establish a presence in a high-tax state, in particular in populous states, where in addition, the fixed costs from operating a warehouse are large. This implicit cost of building a warehouse in a new state provides a unique source of identification since, together with fixed cost differences, tax-inclusive price increases and associated revenue declines serve to offset the benefits from proximity to the consumer.

Thus, in order to examine the role of expansion in increasing concentration, we quantify empirically the trade-off between higher tax-inclusive prices and fixed cost, on the one hand, and additional convenience and shipping cost savings on the other, all while controlling for a potential increase in platform quality through, for example, an expansion of the product assortment. To accomplish this, we combine data on the internet purchase behavior at Amazon and other online competitors by a large representative consumer sample with detailed information on the location and characteristics of Amazon's FCs over time. We take a revealed preference approach similar to that of Holmes (2011), in that we use the fact that entrance of a FC into some states results in lost revenue due to new sales tax liabilities. Therefore, assuming profit maximizing behavior on the part of Amazon implies that entry into a state must also result in either increased willingness-to-pay due to faster deliveries or a reduction of shipping costs. This allows us to form a measure of cost saving by relating the combined change in revenue from tax effects to the reduction in shipping distance through moment inequalities.

The first step in this approach is to estimate a model of demand in order to pin down consumer tax sensitivity and response to shipping times. We specify and estimate demand model using data on both online and offline purchasing of retail goods that we construct by combining the ComScore Web Behavior database, Forrester Research surveys, and household spending data from Environmental Systems Research Institute (ESRI). In the model, a representative consumer in a county chooses her yearly expenditures on four different modes of shopping: (1) Amazon.com, (2) a taxed online competitor (e.g., walmart.com), (3) a non-taxed online competitor (e.g., overstock.com), and (4) a taxed offline competitor (e.g., Wal-Mart). The amount of sales tax charged to the consumer varies both across across time and county due to FC entry and local sales tax laws, implying that

[6] Otherwise, consumers are required to file a use tax return for all of their online purchases. In most cases, consumers do not do this.

variation in spending across these dimensions identifies the consumers' sensitivity to sales tax. Additionally, shipping times vary separately from changes in sale tax due to the variation in the location of consumers relative to the new FC, allowing for identification of the convenience effect. We account for county level variation in spending that is due to consumer demographics and/or the presence of offline retailers using data from the Census. Importantly, we control for changes unobserved platform preferences through a mode-year level fixed effect. This effect accounts for mode-level prices as well as any changes in platform quality which are non-local, such as an increase in product assortment and/or a national improvement in services (e.g., Amazon Prime).

The results of the demand estimation suggest that consumers are sensitive to taxes and that their tax elasticity is around -1.4. This is similar to the estimates from Einav et al. (2014), -1.8, and Baugh et al. (2014), -1.5. The magnitude of this estimate implies that moving from a non-taxed regime to a taxed regime at the average sales tax rate of 6.5%, all else equal, results in reduction of expenditures on Amazon of 9.3%. We find little evidence that the entrance of a FC leads to an increase in demand for Amazon from increased convenience. The lack of effect of shipping time on demand is likely due to the fact that Amazon's local shipping options and prices did not drastically change between 2006-2013, with any national level changes being accounted for through the mode-year fixed effect. Importantly, together these estimates imply that there are overall negative revenue effects of opening a fulfillment center in a new state, meaning this loss must be offset by a decrease in costs. Not surprisingly we find that the fixed effect for Amazon increases steadily from 2006-2013 and is relatively constant for the other two modes, suggesting that preferences for Amazon have increased over this period through price reductions and/or improvements in platform quality.

In the second step, we quantify the shipping cost savings from the larger FC network that offset revenue declines by specifying a profit function where the variable cost of shipping goods to consumers in a county is proportional to the revenue value of the shipment and the shipping distance from the FC to that county. We assume that the observed placement of the fulfillment center network is optimal so that the profit under this structure must be at least as large as under any perturbation of the network. A perturbed network configuration affects revenue through the potentially differing sales taxes, fixed costs through the wages and rents paid in the counties where Amazon FCs are located, and variable costs through changes in the total shipping distance across customers. We compute perturbations to the observed network configuration similar to those of Holmes (2011): we swap the observed opening dates of two FCs.

Calculating profit differences between this set of perturbations and the observed configuration allows us to form moment inequalities from which we estimate bounds of the effect of shipping distance on variable cost using the structure of Pakes et al. (2015) and the methods of Andrews and Soares (2010). Intuitively, alternate network configurations that lead to higher revenue due to sales tax (i.e., they are in a low-populated, low-tax state) but longer shipping routes are informative of the lowest value of the cost-saving-per-mile such that the cost savings from being close to the populated areas outweigh the lost revenue from the tax effect. Similarly, the configurations that lead to lower revenue and shorter shipping routes are informative of the highest value of cost savings such that the increase in shipping cost doesn't outweigh the increase in revenue. We define instruments similar to those in Holmes (2011), adjusted for the nature of our problem where any measurement error in revenue appears not only in the difference in pre-shipping profit, but also the difference in revenue-weighted shipping distance and thus shipping cost.

Estimates imply that it costs Amazon between 58 cents and $1.55 to ship $100 of goods 100 miles, net of shipping revenue. We use these estimates to perform three exercises. First, we quantify the effects of expansion on revenue, fixed costs, and shipping costs from 2006-2018. We find that expansion has led to a $9.6 billion ($2.2 billion) decrease in revenue (profit net of shipping cost) due to the tax effect and a $3.3 billion increase in the fixed costs of operating the network of FCs. However, Amazon has saved between $5 billion and $13 billion in shipping costs by decentralizing the network, implying a net profit increase of between 500 million and $8 billion over this time period. Second, we estimate a measure of the long-run effects of expansion: the change in the 2018 profit margin. We find that the margin is between 5 and 14% higher than what it would have been under the 2006 network, resulting in between $1.6 and $4.5 billion in additional profit in 2018 alone. Finally, our estimates imply that the net average cost of shipping an order, including both the fixed cost of operating the network and the variable shipping cost, has decreased from between $2.69 and $4.67 in 2006 to between $1.09 and $2.14 in 2013.

In order to provide a connection between this cost saving and the increase in concentration, we decompose the estimated preferences for Amazon into a price component and a quality component. We find that, while the relative quality on Amazon is increasing, the relative prices on Amazon are decreasing by approximately 45%, a decrease which mirrors the decrease in shipping cost (55%). While our model cannot directly translate the estimated cost saving to this decrease in prices, this provides suggestive evidence that Amazon's expansion has, at least partially, lead to a competitive advantage through a more efficient network. In order to investigate another possible source of concentration, we correlate the increase in the quality component of the estimated Amazon preferences with an increase in the variety of goods that the platform sells, which many cite as en explanation for Amazon's success.[7] Again, while our model cannot directly speak to this, we show that the assortment on Amazon has increased substantially relative to it's competitors over this time period and that this increase mirrors the increase in the estimated quality of the platform.

This paper is related to several strands of the literature. An increasing body of work focuses on the estimation of demand in online retail markets, with the studies focusing on the effects of sales tax and the gains from variety being the most relevant to our study. Einav et al. (2014) estimate the response to sales tax using eBay data, exploiting the fact that a buyer has the option to buy from an out-of-state seller who does not charge sales tax. They are also able to estimate a

[7]See http://www.nytimes.com/2009/09/20/business/20amazon.html

consumer's sensitivity to distance because they observe the locations of both the buyers and the sellers. Baugh et al. (2014) uses a differences-in-differences approach to estimate the effect of the 'Amazon tax', or changes in the laws that forced Amazon to charge sales tax in various states during 2013-2015. Neither of these papers fully consider substitution to other taxed or non-taxed online outlets. We thus contribute to this literature by expanding our analysis of the tax sensitivity and the estimation of demand for retail goods beyond a single online firm to a large number of online and offline firms.

Papers examining the gains from variety in e-commerce include Brynjolfsson et al. (2003) and Quan and Williams (2016). Brynjolfsson et al. (2003) uses data from online bookstores to demonstrate that consumer welfare gains from increased product variety far outweigh the also sizable gains due to lower prices and increased competition. Findings in Quan and Williams (2016) suggest that the extent of such gains varies significantly across local markets, depending on the availability of competing, localized, assortment in brick-and-mortar stores. In our analysis, on the other hand, we examine the connection between increases in product assortment and increases in concentration. Consumer valuation of product variety provides a large online retailer such as Amazon with a clear advantage over their smaller competitors, which is something that we capture in our demand estimation.

Additionally, to our knowledge, we are the first to examine the possible convenience effects associated with a broader distribution network that cuts down on delivery times, which consumers may value independently of possible tax implications that come with Amazon's FC entry into a state. Our finding that the tax effects dominate any willingness-to-pay response due to increased delivery convenience rules out such demand-side benefits from proximity as a consideration for Amazon's network formation.

Our analysis is also related to recent economics literature examining the interplay of brick-and-mortar retail store locations choices and the retailer's distribution network. Zheng (2014) relates the proximity of a rival's fulfillment center to the chain's expected future entry, while Holmes (2011) estimates the savings in distribution costs associated with clustering stores near a fulfillment center. Both of these studies take the placement of the distribution network as given and use the variation in the network to identify the parameters of interest, whereas we study the network of FCs as a strategic choice variable. Apart from a small operations research literature that looks at the management of distribution networks for online firms (see Agatz et al. (2008) for an overview of this literature), little work to date has studied such classic industrial organization questions in the context of both distribution and online markets.

The remainder of the paper is organized as follows. The next section discusses the background of the sales tax laws and Amazon's FC network. Section 3 introduces the main data sources and Section 4 presents the demand side analysis and results. Section 5 provides estimates of distribution cost savings and Section 6 concludes.

2 Amazon's FC Network and Sales Tax

We obtain information about Amazon's fulfillment centers from the supply-chain consulting company MWPVL, International.[8] MWPVL provides information on the location, size, opening date, and closing date of each FC. Also observable for a subset of the locations is the fulfillment center 'type'. The type of FC is usually defined by the size of the item being shipped and/or the speed of delivery. The primary types are centers that focus on large items that cannot be sent combined with any other products (non-sortable), small items that can be combined in one package (small-sortable) and large items that can be combined into one package (large sortable). In the case of Amazon, other types of distribution centers include Amazon Fresh FCs, which supply Amazon's grocery delivery service, return centers, redistribution centers for third-party distribution, and centers for select specialty items such as jewelry. Starting in 2014 Amazon started to build 'sortation' centers, which are used to sort packages by zip code after they have shipped from a FC, and 'Prime Now Hubs', which handle same-day delivery for the local markets. For the majority of the analysis, we focus on the three primary types of FCs and the Prime Now Hubs, as these are the centers that ship non-grocery items directly to consumers.

For these types, Amazon has expanded from 8 FCs in 2006 to 54 by the end of 2015 and has plans to expand further to over 100 FCs by 2018. Table 1 and Figure 2 demonstrate this expansion. The darker shades in Figure 2 indicate a higher the population and the weighted average sales tax in 2013 is shown. There are a few things to note to about the expansion until 2010. First, Amazon placed FCs in relatively low population states that were close to highly populated areas. For example by 2010, Amazon had two FCs in Nevada, both of which were on the California border close to that state's major cities. Second, they also placed FCs in states with relatively low sales tax. For example, there were FCs in New Hampshire and Delaware, which are both close to major East Coast cities and have zero sales tax. This reflects that sales tax rates are positively correlated with population (across states, correlation of between 0.35 and 0.4 across years), so that entry into a small state near a large state has tax implications for only a small population, while allowing the firm to serve both states' populations more efficiently. Third, when Amazon did expand to highly populated states, they focused on states with a relatively low sales tax rates (e.g., Pennsylvania with an average tax rate of 6.3% in 2014, compared to 6.94% across the top 20 US states in terms of population). For comparison purposes, we include a map of Wal-Mart's distribution centers, which are spread across the country more evenly than Amazon's FCs, at least in the initial stages of Amazon's distribution network.

Given these patterns, it is clear that the early strategy of this expansion depended on tax laws that allow e-commerce firms to avoid charging their customers sales tax, which can be significant, reaching close to 10% in some states. In 1992, the United States Supreme court ruled that the Commerce Clause in the US Constitution "prohibits a State from imposing the duty of use tax

[8] `http://www.mwpvl.com/`.

collection and payment upon a seller whose only connection with customers in the State is by common carrier or by mail." Essentially, the law states that an online retailer does not have to charge a consumer sales tax unless it has a physical presence, or a 'nexus', in that consumer's state of residence. It is the duty of consumers to file a 'use-tax' return every year, which includes purchases from out-of-state vendors. However, very few individuals actually comply with this rule.[9] Accordingly, a large and growing literature examining the effect of sales tax on online purchasing has found significant consumer responses to sales tax rates using, for example, variation in tax rates across municipalities (Goolsbee 2000a, 2000b). Both Alm and Melnik (2005) and Ballard and Lee (2007) find small but significant effects of sales tax on the decision of whether or not to shop online, while Scanlan (2007) finds that this sensitivity is heterogeneous across the level of tax-rates. Ellison and Ellison (2009), Smith and Brynjolfsson (2001), Anderson et al. (2010), and Goolsbee et al. (2010) all find that online shoppers are sensitive to sales tax, typically focusing on purchases in a single product category.[10]

For most online firms, a physical presence would come in the form of an office headquarters or a fulfillment center, implying that these firms likely do not have to charge sales in many states. As the popularity of e-commerce has grown, policy makers have started to suggest that these laws may be giving online firms an unfair advantage over their brick-and-mortar competitors. Further, states are likely losing out on millions of dollars of tax-revenue (see Bruce et al. (2009)). Because of their success and growing market share, Amazon.com has become the focus of politicians' complaints about these sales tax laws.

In the late 2000s, states began to introduce legislation that involved expanding the definition of a nexus to include the existence of 'affiliates', or websites that allow retailers, such as Amazon, to advertise on their site. For example, if a blogger based in Illinois has a link to Amazon on her site, then Amazon must charge sales tax to all the residents of Illinois. Not surprisingly, Amazon and other big retailers responded to this by shutting down their affiliate programs in states that passed these laws.[11,12] This certainly suggests that Amazon wanted to avoid charging sales tax to its customers. They even say as much in their 2008 annual report.[13]

However, it is also clear that as Amazon grew in scale, the network of FCs expanded, presumably

[9] Baugh et al. (2014) quote that 0.2% of people living in Rhode Island, 0.3% of people living in California and New Jersey, 7.9% of people living in Vermont, and 9.8% of people living in Maine report filing use tax returns.

[10] There is also evidence on the response of consumers to sales taxes in offline markets, analyzing the substitution of shopping expenditures across state and country borders (Agarwal et al. (2013) and Asplund et al. (2007)) and over time due to tax holidays (Agarwal et al. (2013)). Chetty et al. (2009) estimates how the saliency of tax rates affects consumer demand.

[11] `http://techcrunch.com/2011/06/10/amazon-shuts-down-associates-affiliate-program-in-connecticut-over-online-sales-tax/`.

[12] `http://www.kansascity.com/news/local/article325412/Amazon-shuts-down-Missouri-associates-program-over-sales-tax-dispute.html`.

[13] On page 16 it states: "A successful assertion by one or more states or foreign countries that we should collect sales or other taxes on the sale of merchandise or services could result in substantial tax liabilities for past sales, decrease our ability to compete with traditional retailers, and otherwise harm our business."

to be closer to population hubs despite sales tax implications and higher fixed costs of warehousing in densely populated areas. For example, by 2014, we see entry into highly populated states such as California and Virginia and high tax states such as Tennessee (9.4% tax rate). Finally, by 2018, Amazon plans to have operating FCs in Illinois, Georgia, Ohio and North Carolina. Table 1 displays the number of FCs, the numbers of states that have a FC, the number of taxed states and associated counties by year, and the average distance between the closest FC to a county, weighted by county's population. Again, we see that as the network expands, Amazon is entering new states, being taxed in a greater number of states, and reducing the distance between the FCs and the consumers. Overall, the pattern of expansion, both in location and in the dynamics, imply that there exists a trade-off between being close to customers and charging them sales tax. Our main goal is to empirically investigate this trade-off.

Table 1 also indicates a strategy of increasing capacity (size) of each FCs between 2006 and 2014, but then decreasing it thereafter. This is primarily due to the building of smaller sortation centers and Prime Now hubs outside of larger metropolitan areas in the latter end of the sample. We ignore sortation centers in our analysis but, in general, we capture the trade-off between the size of an FC and its location through the increasing fixed costs associated with either building a bigger FC or building an FC in a location with high rents and wages.

It is important to note that the entry of a fulfillment center does not necessarily mean that Amazon charges sales tax immediately. For example, Amazon first built a FC in Pennsylvania in 2006, but did not begin to charge sales tax until 2011. This is often due to legal battles with the state government as to what constitutes a nexus. On the other hand, sometimes Amazon charges sales tax even when they do not have a FC in that state. This can be due to changes in state laws (e.g., New York) or because of legal agreements with the state to begin charging sales tax ahead of entry (e.g., Connecticut). The latter explain the, at times, significant discrepancy between the number of states where Amazon's customers pay sales tax and the number of states where Amazon has a FC. Given this, we make various assumptions on how Amazon perceives the relationship between entry and sales tax when making its network location decisions.[14]

[14]In addition to negotiating with states over the effectiveness of its tax collection obligation, Amazon at times receives other forms of government financial assistance when opening a FC in a new locality. The online government subsidy data base Goodjobfirst.org lists 39 different economic development programs that Amazon benefitted from during 2006 and 2013. The vast majority consist of ongoing tax credits and training cost reimbursements, which affect the firm's variable profit. Six consist of grants or reduced interest loans. Since we do not observe Amazon's ultimate take-up of such measures, we incorporate their effect on Amazon's choice of when to enter a location via a measurement error component to variable profit.

3 Data

Consumer Purchases

The primary data source for the estimation of the demand model is the comScore Web Behavior Database. ComScore tracks the online purchasing and browsing activity of a sample of between 50,000 and 100,000 internet users (households) per year. The users give comScore explicit permission to monitor their activity. There are two primary databases, one that records each browsing session regardless of whether or not a purchase was made, and another that records transactions. Because our focus is on buying behavior, we focus on the latter.[15]

For each transaction, we observe a unique household identifier, the time of the purchase, the product category and price for each individual item in the basket, the name of the domain where the transaction occurred, and a 'basket total', which is the total price of the transaction including shipping and taxes. In addition, we observe demographic characteristics for each household such as income, age of head of household, and racial background.[16] The first two columns of Table 2 displays information about the reach of the sample in each year. The sample has shrunk over the years from 86,000 households in 2006 to just 46,000 in 2013. Nevertheless, at least 78% of US counties and every US state plus the District of Columbia (minus Hawaii and Alaska) are represented in the data each year.

We find the comScore sample of households to be generally representative of the Unites States population according the 2010 census, with three exceptions: (1) the head of the household is younger, (2) a higher percentage of the households are white and (3) the household income is higher. All of these facts are likely because the sample is drawn from internet users, who are not perfectly representative of the US population. De los Santos et al. (2012) compare the sample of comScore users in 2002 and 2004 to the Computer Use Supplement of the Current Population Survey and find that the sample generally compares well with the population of online shoppers. However, to account for the possibility that comScore may be over or under sampling certain demographic groups, we adjust our data using sampling weights from the census. We bin each comScore household into categories based on income, age, and racial makeup, and calculate relative sampling weights based on the prevalence of each category in the comScore data relative to its prevalence in the 2010 census at the county level.

Table 2 also provides a first look at the purchasing patterns in the comScore sample. The third and fourth columns display the yearly average online expenditures and transactions per household, and the fifth column shows the percentage of households with zero transactions. Note that we have limited the data to only transactions in product categories that Amazon sells. Therefore,

[15] ComScore is an internet analytics firm that provides data to Fortune 500 companies and large media organizations. See De los Santos et al. (2012) for a deeper discussion of comScore's services.

[16] The data contain a single variable indicating a range of household income. To supplement this information, we create a continuous measure of the expected income for each county within a demographic group using information on county level income distributions in Census data.

we omit purchases in categories such as food, travel, and online dating. We further adjust yearly expenditures for the number of weeks we observe the household in browsing data.[17]

Interestingly, we see decreasing average expenditures and transactions over time. Column 5 indicates that this is mostly due to an increasing percentage of households with zero online purchases. This is not in line with anecdotal evidence about the take-up of online shopping.[18] It is possible that the sampling procedure is the cause for this, but a more plausible explanation is that comScore is not recording the transactions for some households. This could be due to the household deactivating the comScore behavior monitor or using a second computer (e.g., at work) or a mobile device for their online shopping. Because we do not know which households truly have zero expenditures, rather than using other devices for purchases, we choose to exclude all households with zero transactions from the sample. In order to account for the extensive margin, we supplement the comScore data with survey data from Forrester Research.

Forrester Research conducts annual surveys of the online shopping behavior of a representative sample of the US population, including whether or not the respondents made any online purchases in the last three months, with the survey being conducted between March and July, depending on the year.[19] In the final column on Table 2, we see that there is an increasing number of households shopping online, according to the survey. We match the Forrester and comScore data based on demographic groupings, and calculate expected expenditures for a household in the comScore sample that accounts for the propensity of not making any online purchases in their demographic category in the Forrester data.[20] See Appendix A for details. After adjusting the comScore data with Forrester's extensive margin, we see a pattern that is more in line with the anecdotal evidence: expenditures and transactions are both increasing over time. Figure 1 illustrates Amazon's growing market share in expenditures over time, growing from under 10% in 2006 to 25% in 2013, which is in line reported estimates.[21]

One limitation of the comScore data is that they do not record Amazon Marketplace purchases separately from Amazon purchases. A worry may be that consumers find a tax-free seller on Amazon Marketplace in a state in which Amazon itself must charge sales tax. This implies that we may underestimate the effect of sales tax because we would attribute a lack of response in states where Amazon charges taxes to a low sensitivity, when it is actually coming from consumers purchasing from a non-taxed vendor on Amazon Marketplace. However, most laws dictate that even marketplace vendors have to charge sales tax to customers in the states where their goods are

[17]We calculate the average weekly expenditures over the number of weeks we observe the household with any browsing activity and then multiply this by 52 weeks.

[18]`http://www.statista.com/statistics/183755/number-of-us-internet-shoppers-since-2009/`.

[19]Note that the survey was only available for 2006-2007 and 2010-2013. We interpolate linearly for the intervening years to construct predicted propensities of buying online.

[20]Similar to the ComScore data, the income variable in Forrester is binned. Therefore, we create a continuous measure of expected income in the same manner.

[21]Internetretailer.com reports a 23% share in the 2nd quarter of 2014 `https://www.internetretailer.com/2014/10/23/amazon-q3-revenue-increases-20`

housed.[22] A recent survey of sellers indicated that 79% of all sellers and 88% of high volume sellers (+$1 million in sales) use the Fulfilled by Amazon (FBA) program, meaning that the sales tax implications of the network are likely the same for many of these vendors.[23] However, we perform a robustness check in Appendix C of the demand model that accounts for this possibility.

Finally, in order to construct the amount of household spending at all online and offline retailers for the types of goods in our sample, we use a combination of data from ESRI and SimplyMap. The former provides a measure of the county level household expenditures on all retail, including food, while the latter provided the county level household expenditures on food. Because we exclude food from our online sample, we subtract the food expenditures from the total to calculate our measure of total spending.

Amazon Fulfillment Center Network

We acquire information about each of the FCs from MWPVL, International. From the opening date and the location of the FC, we calculate the straight-line distance between each FC's street address and the population weighted centroid of every county in the US, and take the minimum of this by county and year. The distance serves as a measure of how far Amazon must ship the good between its FCs and the consumers. It is apparent from Figure 2 that Amazon often opens FCs that are close to an existing FC, which may be to increase capacity or to have FCs of different types near one another. Because of this, we define a FC cluster as a group of FCs that are in the same state and within 20 miles of each other. In the analysis below, we assume that Amazon decides where and when to build a **cluster** of FCs rather than each individual FC in the cluster, and that the observed opening date and shipping distance are associated with the first FC built in the cluster. Note that this does not affect the demand analysis because the tax implications are associated with the first FC built with or without clustering.

Because straight-line distance may not be a perfect measure of shipping speeds, we also obtain the US Postal Service's shipping times between each three digit US zip code and find the minimum shipping time between each county and each FC three-digit zip code.[24] The shipping times are for four different classes of mail: first, priority, standard, and package. These classes differ slightly in the size and number of packages allowed in the delivery. For a detailed description, see usps.com.[25] We were unable to find shipping times for the other shipping services used by Amazon, UPS and FedEx, but they are likely very similar to USPS.

MWPVL also provides information about the size of every FC in square feet and the number of employees for about half of them. To fill in the missing information about the number of employees,

[22] `http://www.avalara.com/learn/whitepapers/fba-sellers-guide-sales-tax/`.

[23] http://www.webretailer.com/lean-commerce/amazon-sellers-survey-2016/

[24] We choose one three digit zip code per county after verifying that shipping times do not vary much within a given county.

[25] `http://pe.usps.com/businessmail101/classes/welcome.htm`.

we assume that the number of employees per square foot is the same for FCs of the same type. The last two columns of Table 1 provide the average number of employees and size for FCs which are open in a given year. There is a general pattern of an increase in size (and employees) early in the sample period, followed by a decrease as Amazon began to build the smaller Prime Now hubs starting in 2014. Note that we do not observe any information about the FCs for Amazon's competitors, so we cannot speak to competition among online firms in terms of their FC network choices.

Finally, we obtain information on the extent to which a given fulfillment center can be used to satisfy same day shipping orders during our sample period. Specifically, we found the date of implementation for Amazon's early version of same day shipping, called 'Local Express Delivery', from various news sources. Amazon began this service, for which customers are charged a fee, as early as October, 2009.

Other Sources of Data

In addition to the above sources, we observe state, county and local sales tax rates from Tax Data Systems, now part of Thomas Reuters. For each year and county we calculate the average tax rate, as local taxes can vary within a county and may change mid year. The overall average tax rate is around 6.5%. There is variation in sales tax rates across all counties, with a standard deviation of about 1.6% in every year, and variation across counties within a state, with an average standard deviation of about 0.3%. In addition, tax rates vary across time, as between 30 and 65% of counties experienced tax rate changes each year and nearly 78% of the counties experienced at least one tax rate change over our sample period.

We collected county-level wage information from the BLS's Quarterly Census of Employment and Wages and used the annual wage of a 'Retail Trade' employee. While this is not the exact industry in which the Amazon FC employees work, we believe the wages are likely to be similar in a given county. This is evidenced by the fact that the overall average hourly wage of a 'Retail Trade' employee is close to that of an Amazon Associate as reported by glassdoor.com. The land rents were computed using data from SNL Financial's real estate research. SNL Financial is a subsidiary of S&P Global Incorporated and it provides detailed real estate data for most publicly traded U.S. Real Estate Investment Trusts. From SNL, we observe the rent paid and the size of the properties classified as "Warehouses or Distribution Centers". Using these, we create a rental rate per-foot measure for all the counties appearing in the SNL data. We note that there are some counties that have an Amazon FC that do not appear in the SNL data (26 out of 60). For these counties, we create a state-level average ratio of rental rates to residential property values, where residential property values were obtained from the United State Census.[26] We then calculate the rental rate as the property value in the missing counties times this ratio. Finally, we collected county level

[26]The data was compiled by `http://metrocosm.com/get-the-data/`

demographics from the 2010 Census, along with annual information on the number of small and large brick-and-mortar retailers from County Business Patterns.

Finally, we use Amazon's annual reports to provide additional information about Amazon's finances. That is, for each year, we obtain the aggregate sales in relevant categories, the cost of goods sold, which includes inbound and outbound shipping costs, and the aggregate outbound shipping cost.[27] This allows us to formulate the gross profit margin for Amazon, net of outbound shipping cost, which we discuss in more detail in Section 5.

4 Demand Analysis

We begin the demand analysis by providing reduced form evidence that consumers are sensitive to sales tax and that their are no convenience effects of expansion. First, we estimate the transaction level effect of taxes on the likelihood a consumer purchases from Amazon. That is, we run the following linear probability regression:

$$Pr(A_{hijt} = 1) = \beta_0 + \sigma \ln(1 + \tau_{it}\mathbf{1}_{it}^{taxable}) + \gamma d_{it} + \lambda_i + \lambda_j + \lambda_t + \epsilon_{hijt}$$

where each observation is a purchase occasion h, from a consumer in county i, in year t, for a product in category j. The dependent variable is a dummy variable equaling one if the consumer purchases the product from Amazon.com and zero otherwise. The tax rate is given by τ_{it}; it varies depending on the county tax rate and whether or not Amazon has to charge sales tax to customers in that county in a given year. The measure of the shipping speed from Amazon to the consumer is given by d_{jt}, which varies depending on the specification. We do not include the item's price or the item's price at Amazon relative to the price charged by an outside seller because we observe the price only on the platform where the transaction occurs, but not on alternative platforms.[28] However, we include product category and time fixed effects to account for average price differences across online retailers at the category level and across time.

The result in the first column of Table 3, which excludes any measure of shipping speed, indicates that consumers are sensitive to taxes and that they reduce their purchases from Amazon by around 14% for a one percentage point increase in sales tax, result which is statistically significant at the 10% level.[29] In columns 2 through 7, we include various measures of shipping speeds to account

[27] `http://phx.corporate-ir.net/phoenix.zhtml?c=97664&p=irol-reportsannual`.

[28] Nevertheless, including a variable that equals to Amazon's price for Amazon purchases and zero otherwise results in similar estimates.

[29] In a few states (i.e., MN, NJ, PA, RI, and VT), retailers are not required to charge sales tax on clothing. There is an "Apparel" category in the comScore data, but we believe this may not accurately define the goods which are not taxed for a couple of reasons. One, products reported in the Apparel category come from a wide range of goods, some of which are not clothing (i.e., accessories). Two, upon spot check, products are often mis-categorized. Finally, there are a number of exceptions of these tax laws, which we cannot systematically identify. For example, sporting wear and formal wear are not tax exempt in Pennsylvania. Despite this, we perform the two robustness checks to this analysis. First, we remove the states which are exempt and second, we remove any purchases which fall into

for the fact that consumers may value faster shipping times from Amazon. The measures are whether or not the county had access to "local express delivery"', Amazon's early form of same day delivery, the log of the distance between the centroid of the consumer's county and the closest FC, and the shipping times from USPS for the four different classes of mail: priority, first, standard, and package. Since the variation in shipping days is limited, we aggregate shipping times into two categories, fast and slow shipping speed between two three-digit zip codes, excluding the slow shipping speed indicator from our specifications. We find that the tax effect is robust to including the variety of proxies for shipping time. Neither the local express delivery dummy nor the log of the shipping distance have a statistically significant effect on the propensity to purchase from Amazon. This is also the case for all mail classes, where the probability of purchasing from Amazon does not change by a significant amount compared to customers in zip code locations with a slower shipping speeds, all else equal.

There are a number of explanations for the results on the shipping speed coefficients. It could be that our measures do not accurately reflect shipping times. Another explanation is that consumers simply do not have tastes for faster shipping, or have very nonlinear preferences for shipping times given Amazon's recent expansion to same-day shipping. Finally, it could be that the expansion of the network of fulfillment centers did not result in faster shipping times to local consumers, or if it did, it was a small number of consumers who were affected. We believe this to be the most plausible explanation.

Next, we perform a differences-in-differences analysis of the response in Amazon expenditures to a change in the tax status of a county. We summarize a county's tax status in an indicator variable, $\mathbf{1}_{it}^{taxable}$, that is one if a purchase from Amazon.com by a household in county i in year t is subject to sales tax. The tax status variable changes over time due to the entry of a FC and/or due to the future entry of a FC where the state insisted on collecting sales taxes immediately following the initial agreement. The regression model is given by:

$$e_{it} = \beta_0 + \sigma \mathbf{1}_{it}^{taxable} + \gamma d_{it} + +\lambda_t + \lambda_i + \epsilon_{it}$$

where e_{it} is the log of aggregate expenditures on Amazon from county i in year t. Aggregate expenditures are the sum of expenditures on Amazon in year t for the comScore households who live in county i. Results in Table 4 suggest that a change in the tax status of a county results in somewhere between a 9.5 and 10.6% reduction in expenditures on Amazon, with this effect being significant at the 10% level in all specifications and at the 5% level for specifications (5) through (7). This exercise is essentially the same as the one performed in Baugh et al. (2014) with different data, and they estimate a tax effect of about 9.5%. We once again include measures of shipping speeds in this demand estimation and find little evidence that they significantly affect the

the Apparel category. We find that the tax coefficient is significant at the 10% level and of similar magnitude as the presented results.

purchasing behavior of consumers.

A weakness shared by both of the above exercises is that we do not consider the substitution between Amazon and other taxed and non-taxed shopping options. Because of this, any reduction in Amazon's transactions due to higher taxes does not equate directly to a tax elasticity since these transactions may be substituted towards other taxed outlets (e.g., walmart.com or Wal-Mart). Therefore, we specify and estimate a model of demand for retail goods across all modes of online and offline shopping in the following sections that allows us to determine the effect on overall transactions of higher sales taxes.

Empirical Model

We specify a model of demand where a consumer chooses how much money to spend on retail goods in each of four different modes of shopping. The modes are Amazon, $j = 1$, taxed online competitors, $j = 2$, non-taxed online competitors, $j = 3$, and offline competitors, $j = 0$.

We follow Einav et al. (2014) and specify a CES utility function. Specifically, a representative consumer from county i solves the following problem in year t:

$$\max_{q_{i0t},..q_{i3t}} \left(\sum_{j=0}^{3} \left(\frac{q_{ijt}}{\zeta_{ijt}} \right)^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}$$

$$s.t. \sum_{j=0}^{3} p_{ijt} q_{ijt} \leq w_{it}$$

where q_{ijt} represents the quantity purchased via shopping mode j in time period t, p_{ijt} is the price of purchasing one unit of goods via shopping mode j, ζ_{ijt} is the taste for mode j, and w_{it} is household i's budget for retail goods. The elasticity of substitution between the four modes is given by σ. Solving for the optimal amount of expenditures in each online mode results in:

$$e_{ijt} = \frac{(p_{ijt}\zeta_{ijt})^{1-\sigma}}{P_{it}^{1-\sigma}} w_{it},$$

where P_{it} denotes a weighted-average price index across all four modes. Dividing this by the expenditures for $j = 0$ and taking logs gives:

$$\ln(e_{ijt}) - \ln(e_{i0t}) = (1-\sigma)(\ln(p_{ijt}) - \ln(p_{i0})) + (1-\sigma)(\ln(\zeta_{ijt}) - \ln(\zeta_{i0t})) \quad (1)$$

where the price of each mode can be written as:

$$p_{ijt} = (1 + \tau_{it} \mathbf{1}_{ijt}^{taxable}) \tilde{p}_{ijt}$$

The tax rate in county i and year t is given by τ_{it}, meaning $\tilde{p}_{ijt}$ is the tax-exclusive price of buying goods, which may include shipping charges. Note that the tax liability may vary across j as the non-taxed competitor never charges sales tax and Amazon does not charge sales tax in a number of states and years. We indicate whether mode j entails a sales tax liability for the customer at time t through the indicator $\mathbf{1}_{ijt}^{taxable}$; purchasing from an offline competitor always implies that the customer has to pay sales tax. With this, equation 1 becomes:

$$\begin{aligned}\ln(e_{ijt}) - \ln(e_{i0t}) =& (1-\sigma)(\ln(1+\tau_{it}\mathbf{1}_{ijt}^{taxable}) - \ln(1+\tau_{it})) \\ &+ (1-\sigma)(\ln(\tilde{p}_{ijt}) - \ln(\tilde{p}_{i0t})) \\ &+ (1-\sigma)(\ln(\zeta_{ijt}) - \ln(\zeta_{i0t}))\end{aligned} \tag{2}$$

We model the taste for the each mode $j \in \{1, 2, 3\}$ as:

$$\zeta_{ijt} = \exp(\xi_{jt} + \gamma d_{it}\mathbf{1}_j^{Amazon} + \beta_j Z_{it} + \lambda_j C_{it} + \epsilon_{ijt})^{\frac{1}{1-\sigma}}$$

and assume that the preference for the offline shopping mode equals:

$$\zeta_{i0t} = \exp(\Delta\xi_{i0t} + \epsilon_{i0t})^{\frac{1}{1-\sigma}}$$

The term $\Delta\xi_{i0t}$ is the sum of the county level preferences for online shopping, ξ_{i0}, and the time varying preferences for online shopping within the county's Census District, $\Delta\xi_{r0t}$.

The mode-year effect, ξ_{jt}, accounts for the time varying preferences for shopping mode j and can be thought of as the mode's quality that doesn't vary across locations. Aspects such as product variety, return policy, and customer service likely play a major role in determining ξ_{jt}. For Amazon specifically, things such as availability and take-up of the Prime service, which all customers have access to regardless of their location, would be included in ξ_{jt}.

County level variation in preferences comes from Z_{it}, which is a vector of county level demographics, C_{it}, which is a vector of variables measuring the level of offline competition, and d_{it}, which is the measure of shipping speed from Amazon. Notice that we allow the effects of demographics and competition to vary across shopping modes. For example, having a brick-and-mortar Best Buy in a county may affect the purchases through taxed online competitors (including bestbuy.com) differently than purchases from Amazon.

The demand shocks are at the county, time, and mode level, and are assumed to be iid across these delineations. Under these assumptions, equation 2 becomes:

$$\begin{aligned}\ln(e_{ijt}) - \ln(e_{i0t}) =& (1-\sigma)(\ln(1+\tau_{it}\mathbf{1}_{ijt}^{taxable}) - \ln(1+\tau_{it})) \\ &+ \gamma d_{it}\mathbf{1}_j^{Amazon} + (1-\sigma)(\ln(\tilde{p}_{ijt}) - \ln(\tilde{p}_{i0t})) \\ &+ \beta_j Z_{it} + \lambda_j C_{it} + \xi_i^o + \xi_{rt}^o + \xi_{jt} + \epsilon_{ijt} - \epsilon_{i0t}\end{aligned} \tag{3}$$

We make the assumption that the prices for the online shopping modes do not vary across counties, or that $\tilde{p}_{ijt} = \tilde{p}_{jt}$. While some online firms appear to price discriminate based on a consumer's location, there is limited evidence that this is widespread.[30] Further, Amazon attempted to implement price discrimination in 2000, and quickly abolished it after a backlash from customers.[31,32] Finally, we assume that the base price of the offline shopping option remains constant over time (i.e., $\tilde{p}_{i0t} = \tilde{p}_{i0}$) and that any variation can be captured through C_{it}. With this, we can re-write equation 3:

$$\begin{aligned}
\underbrace{\ln(e_{ijt}) - \ln(e_{i0t})}_{\tilde{e}_{ijt}} =& (1-\sigma)\underbrace{\left(\ln(1+\tau_{it}\mathbf{1}_{ijt}^{taxable}) - \ln(1+\tau_{it})\right)}_{\text{Price variation } (\tilde{\tau}_{ijt})} \\
&+ \gamma d_{it}\mathbf{1}_j^{Amazon} + \beta_j Z_{it} + \lambda_j C_{it} + \underbrace{\xi_{jt} + (1-\sigma)\ln(\tilde{p}_{jt})}_{\text{Mode-year FE}} \\
&- \underbrace{\Delta\xi_{r0t}}_{\text{Region-year FE}} - \underbrace{\bar{\xi}_{i0} + (1-\sigma)\ln(\tilde{p}_{i0})}_{\text{County FE}} + \underbrace{\epsilon_{ijt} - \epsilon_{i0t}}_{\text{Residual}}
\end{aligned} \tag{4}$$

This equation says that the difference in expenditures between mode j and the offline mode is a function of the difference in sales tax between mode j and the offline mode, the convenience effect, a time varying mode level effect that includes the price level of mode j, the relative effect of demographics and local competition on mode j, a region-specific time trend, a county effect, and an iid demand shock.

County Level Expenditures

To form the expenditures for the representative consumer ($\tilde{e}$), we calculate the population weighted yearly average expenditures on each shopping mode for each county. We define a taxed online competitor as one that has a large offline presence such as gap.com, walmart.com and target.com and a non-taxed competitor as one without a national offline presence, such as overstock.com. We note that neither of these groups are perfectly defined, as some sites that we classify as having a national offline presence may not have a store in every state and sites that we classify as not having an offline presence surely have a headquarters and/or a fulfillment center located in at least one state. Table 5 displays the top ten online stores in each of the taxed and non-taxed categories. Overall, we classify about 34% of the websites that compete in product categories that Amazon carries as taxed competitors. We also exclude any expenditures in product categories that Amazon does not cover and exclude any websites that sell only in these categories. Examples of excluded sites are dating websites (e.g., Match.com), travel websites (e.g., orbtiz.com) and food delivery sites (e.g., dominos.com).

[30] http://www.wsj.com/articles/SB10001424127887323777204578189391813881534.
[31] http://www.bizjournals.com/seattle/stories/2000/09/25/daily21.html.
[32] http://news.cnet.com/2100-1017-240700.html.

Under these restrictions, we calculate the yearly expenditures for the three online shopping modes for each individual household. The level of household expenditures is then adjusted using the Forrester survey data and the average expenditures for a county are calculated using demographic weights from the 2010 census. Finally, expenditures for the offline mode are calculated by subtracting online expenditures from the total amount of expenditures on retail goods calculated from the ESRI and Simplymap data. See Appendix A for further details on how expenditures are calculated.

Estimation

We estimate equation 4 using OLS with county, mode/year, and region/time level fixed effects. Included in C_{it} is the total number of retail establishments and the number of establishments with more than 50 employees (i.e., large retailers) in county i in year t to account for the fact that large offline retailers (e.g., Wal-Mart) may have a different impact on online demand than small local retailers. County level demographic included in Z_{it} are median income and the distribution of ethnicity.

Identification of the tax sensitivity parameter comes from two primary sources of variation. First, there is variation in tax rates in a county across time due to changes in local laws and/or entry of an Amazon fulfillment center. Therefore, changes in expenditures between taxed and non-taxed modes as a result of these changes helps to identify σ. This is similar to the traditional identification argument in difference-in-differences models. Second, there is variation of tax rates across counties within a state, which allows us to utilize the covariation between level of the response to changes in the tax status and the tax rate.

Given this level of variation, we make the reasonable assumption that the changes in tax rates due to either Amazon's expansion and/or local laws are exogenous to the unobserved local demand shocks. This reflects that Amazon's entry decisions are made at a higher geographic level than the county and not made purely in response to a local annual demand shock in a specific county. To investigate the sensitivity of our results to this assumption and, similarly, to the assumption above that the price of the offline option does not vary over time, we have also estimated the above specifications with county/year and state/year fixed effects. We find similar results, but choose to proceed with the main specification to avoid identifying the tax effect purely from variation in taxes across modes within a county or state, since we observe only three modes and at most two tax levels across the modes.

As in the reduced form regressions above, we explore whether consumers are sensitive to shipping speeds. We identify this effect using changes in the distance and shipping times resulting from the expansion of the FC network. Specifically, the effect of a new FC on shipping times varies both within and across state borders, as it depends on the location of the consumers. Note that we do not observe the locations and thus distance to Amazon's competitors' fulfillment centers, but we assume that the time-varying preferences, ξ_{jt}, capture any overall effect of changes in shipping times for a given mode. Therefore, the effects of shipping speeds are Amazon specific.

Finally, a nice feature of the data is that entry of a fulfillment center into a state does not always lead to Amazon having to charge sales tax (e.g., the Pennsylvania example from above). Amazon may also have to charge sales tax in a state without having a fulfillment center. Therefore, even within a state, we sometimes see variation in distances without variation in taxes and vice versa.

Matching Total Sales

One of the goals of the demand model is to be able to predict Amazon's total sales under different configurations of the network. However, because of two data limitations, the model presented above may not accurately predict total revenue. First, we are missing sales made through other channels such as mobile devices or second computers at home or work. This issue exists across all three online modes of shopping, suggesting that we would underestimate revenue for all modes. Second, we cannot separate Amazon purchases from Amazon Marketplace purchases, implying that our model predicts the sum of Amazon's own sales and its sales from the sellers on Marketplace. This issue exists only for Amazon and results in an overestimate of its sales.

In order to address this, we supplement the ComScore data with information on Amazon's yearly revenue obtained form their financial reports and information on yearly online sales from the US Department of Commerce.[33, 34] We include two additional yearly parameters, which we call 'multipliers', that are identified by matching the models predictions with these sources. Specifically, we iterate over the OLS regression of equation 4 for different values of α_{tA} and α_{tM}, where the data used for the regression is generated by multiplying each household's expenditures in year t on Amazon by α_{tA} and each household's expenditures in year t on modes 2 and 3 by α_{tM}, until the predictions from the model equal the reported revenue.[35]

The implicit assumption in this procedure is that the share of missing transactions and Amazon Marketplace expenditures are not a function of the network of FCs. That is, the shares of these expenditures do not vary across counties. Therefore, the primary effect of adjusting expenditures in this way is in the estimates of the mode-year fixed effects and the other parameters remain largely unchanged when estimating the model without the multipliers. Importantly, there is very little change in the estimated tax sensitivity and the convenience effect.

[33] We use the figures of sales from North America in the "Media" and "Electronics and Other General Merchandise" categories. This excludes the "Other" category, which is revenue from "Non-Retail Activity" such as Amazon Web Services (AWS).

[34] `https://www.census.gov/retail/ecommerce/historic_releases.html`

[35] An additional complication is that the model's predictions are in terms of a subset of product categories, whereas the reports from the Department of Commerce are across all categories. Therefore, within this procedure, we divide the predictions of total US online sales by the yearly share of expenditures in the non-missing categories.

Results

Results of 6 different specifications are presented in Table 6, where each specification varies by the measure of shipping speed included. Specification (1) does not include any measure, specification (2) includes whether or not local express delivery is available, specification (3) includes the log of the distance to the shortest FC and specifications (3)-(6) include variables for shipping speed in four different classes of delivery. In the latter, the longest delivery time is the excluded dummy variable. The estimated value of σ, or tax sensitivity, is around 1.43 and is significant at the 5% level in all specifications. These estimates are similar to the ones reported by both Einav et al. (2014) and Baugh et al. (2014).

The estimates imply that going from not charging sales tax to charging the average sales tax of 6.5% would reduce expenditures on the given mode by 9.3 percent. That is, if Amazon agrees to charge sales tax in a state through either an agreement with the state government or because they build a fulfillment center, they can expect their revenue to decrease by up to 9.3%, all else equal. In 2008, New York state passed a law that required Amazon to collect sales tax. Our estimates imply that this reduced Amazon's revenue in New York by around 12.2%.

Once again, we do not find strong evidence that our measures of shipping times drive substitution between Amazon and other shopping outlets. This is further evidence that the expansion of the network did not significantly change the shipping times from Amazon, or in other words, the convenience effect of expanding the network is approximately zero. For further robustness analysis of these demand estimates, see Appendix C.

We estimate values of α_{tM} that are increasing over time, rising from 1.00 in 2006 to 2.96 in 2013, a number which implies that we are missing about two thirds of sales in the final year of our sample.[36] The increasing rate is likely due to the fact that mobile shopping has increased over this time period. In general, the estimates of these multipliers are in line with the reported shares of purchases that online shoppers make via a mobile device (24.6% in 2014) and the incidence of at-work-online shopping (47% in 2014).[37] The values of the Amazon specific multipliers are also increasing over time, from 1.03 in 2006 to about 1.69 in 2013. [38] Assuming that the rate of missing observations is the same for Amazon and modes 2 and 3, the 2013 figure implies that around 57% (1.69/2.96) of the observed revenue is due to purchases directly from Amazon. This is in line with reports that approximately 40% of sales where through third party sellers on Amazon in 2013.[39]

The estimates of the time varying mode fixed effects are presented in Table 7, with the excluded effect being mode 3 in 2006. Amazon's effect is increasing over time, which is not surprising considering the increase in market share over this time period. Both the taxed online competitors'

[36]The multipliers are 1.02, 1.35, 1.36, 1.41, 1.71, 1.81 2.97, and 2.93 for the years 2006-2013.

[37]See `https://www.internetretailer.com/2015/08/18/mobile-commerce-now-30-all-us-e-commerce` and `http://www.careerbuilder.com/share/aboutus/pressreleasesdetail.aspx?sd=12%2F1%2F2014&id=pr854&ed=12%2F31%2F2014`.

[38]The multipliers are 1.03, 1.39, 1.50, 1.55, 1.52, 1.50, 1.01, and 1.69 for the years 2006-2013.

[39]`https://www.statista.com/statistics/259782/third-party-seller-share-of-amazon-platform/`

and non-taxed competitors' effects are also increasing, but at a lower rate compared to Amazon. Finally, the estimates of the parameters on C and Z are presented in Table 8. This table shows a negative effect of offline competition on expenditures for all modes of online shopping, with the effect of large offline competitors (e.g., Wal-Mart or Target) being stronger. However, the effect is only significant for taxed online competitors, implying that online firms with an offline presence face more competition from large offline retailers. A majority of the demographic coefficients are not significant except for income, which indicates that higher income households tend to spend more online. Note that because we do not have the ethnicity variables varying over time, we exclude mode 3's effects.

Using the estimates of specification 1, we predict yearly revenue for each mode from 2006-2013, which are displayed in the upper panel of Table 9. We see a general increase in spending across all modes, with Amazons sales increasing faster than the other two modes. We examine patterns of spending across the shopping modes in Figure 3a, which displays the market share of the online shopping modes over time. Note that, while this is the market share of online purchases only, a figure considering the offline option would feature a similar pattern but with significantly smaller shares. Amazon's market share has steadily increased from under 10% in 2006 to 25% in 2013, while mode 2's share has remained constant around 40% and mode 3's share has decreased from around 50% to 35%. This is concurrent with the expansion of Amazon's FCs documented 1 that resulted in an increase in the number of counties taxed and a decrease in the shipping distance.

Finally, in order to take a first look at the sources of these changes in shares, we construct a measure of the mode preferences:

$$\bar{\xi}_{jt} = \hat{\dot{\xi}}_{jt} + \hat{\beta}_j \bar{Z}_t + \hat{\lambda}_j \bar{C}_t$$

where $\hat{\dot{\xi}}_{jt}$ is the estimated FE and $\bar{Z}$ and $\bar{C}$ are the averages of the demographic and competition variables, and display them in Figure 3b. Not surprisingly, the mode preferences follow a similar pattern as market share, with Amazon's effect increasing significantly compared to the other modes. Due to the way we specified demand, this could be due to a relative decrease in prices and/or a relative increase in platform quality.

To the extent that it is due to quality, features such as an increase in assortment and/or an improvement in services through Amazon Prime likely play a large role in this increase. To the extent that it is due to price reductions, it could be a result of the gain in efficiency from expansion being passed through, at least partially, to the consumer. We explore these issues further in our supply side analysis and subsequent discussion.

Projecting Revenue

To calculate revenues for 2014-2018, we need to obtain future values of the mode specific year effects and the region time effects. We do this by using the growth of these effects from 2006-2013 to predict future effects.[40] The lower panel of Table 9 reports the total predicted revenue from 2014-2018. Because we observe Amazon sales in 2014 and 2015 from their financial report, we can assess how well this procedure does in predicting sales in these years. We find that we over predict revenue by around 5% in both years (53.6 billion versus 50.8 billion in 2014 and 65.9 versus 62.9 in 2015), but that the prediction is reasonably close. We also predict that the mode 2 expenditures rise from \$62 billion in 2014 to \$80 billion in 2018 and mode 3 expenditures rise from \$49 billion to \$61 billion.

5 Quantifying Cost Savings

We now turn to estimating the cost savings of expanding the network of FCs. We posit that these cost savings come the fact that, after expansion, Amazon shortens the outbound shipping distance (i.e., from the FC to the customer) that is handled by one of the contracted shipping companies. That is, with more localized FCs, the outbound shipping distance from the FC to the shipping service's sorting center becomes shorter. Further, the expansion may even remove the leg going from the FC to the sorting facility all together, allowing the shipping service to deliver the package directly from the FC to the customer. Under either scenario, the service's costs are lower, giving Amazon more bargaining power when negotiating their per package shipping rates.

One concern is that the expansion of the network may also lead to longer inbound shipping routes from the supplier to the FC resulting in higher inbound shipping costs that might change Amazon's wholesale price if suppliers could successfully pass on their higher cost. However, the fact that Amazon's suppliers already deliver goods to many other retailers across the US and/or have their own widespread distribution network alleviates this concern to a certain extent. For example, Barnes & Noble requires book publishers to deliver inventory directly to stores, while Walmart – similar to Amazon – requires shipment to distribution centers. With Amazon's expansion of its distribution network, the average distance from a given Amazon fulfillment center to the closest Walmart fulfillment center falls from 92.2 miles in 2006 on average to 65.4 miles in 2013. With economies of scale in distribution across retailers, this should limit potential increases in shipping cost incurred by Amazon's suppliers from serving a more decentralized network.

[40] Specifically, we run the following regression $\hat{m}_t = \gamma_0 + \gamma_1 \ln(t) + \gamma_2 \ln(t)^2 + \epsilon_t$, where $\hat{m}_t$ represents the estimate of the effect from the 2006-2013 data. Using the estimates of γ we predict future values of m_t.

Profit Function

To quantify cost savings, we formulate a profit function for Amazon that depends on the network of FCs it operates in year t, denoted as a_t, in a similar manner as Holmes (2011). In county i, Amazon generates $R_{it}(a_t)$ in revenue, which is a function of a_t due to sales tax implications. Variable profit in year t from county i is given by:

$$\pi_{it}(a_t;\theta) = \mu R_{it}(a_t) - \theta d_{it}(a_t) R_{it}(a_t)$$

where μ is the share of revenue that Amazon receives as profit net of any shipping costs and the variable shipping cost is given by $\theta d_{it}(a_t) R_{it}(a_t)$. We set $\mu = 0.23$ based on information obtained from Amazon's financial reports. See appendix B for details. The parameter θ measures the shipping cost per dollar of goods sold per mile and $d_{it}(a_t)$ is the distance from consumers in county i to the closest FC. The biggest difference with Holmes (2011) is that we allow variable costs to vary across network location decisions, whereas Holmes (2011) assumes that distribution costs are fixed. We assume that this variation in costs comes from differences in the shipping distance from a FC to consumers in county i. Therefore, total profit for Amazon, starting in 2006 is given by:

$$\Pi(a;\theta) = \sum_{t=2006}^{\infty} \beta^{t-2006} \sum_i \mu R_{it}(a_t) - \theta d_{it}(a_t) R_{it}(a_t) - F_{it} a_{ti} - S_t(a_{it} - a_{it-1}) \qquad (5)$$

where a represents the 'rollout' of the entire network, $a_{it} = 1$ if there is a FC in county i in year t, and β is the discount factor, which we set to 0.95. The fixed costs of the operating a FC in county i and year t are given by F_{it}. They are assumed to be made up of wages paid to employees of the FCs in county i, called 'Associates' and denoted by A_{it} and the land rents for the total square footage of building space in county i, L_{it}:

$$F_{it} = w_{it} A_{it} + r_{it} L_{it}$$

Finally, the sunk cost of opening a FC is S_t, which does not vary by county. These fixed costs include machinery and other materials that are needed to open a FC.

Therefore, the profit maximization problem, in terms of network formation, can be written as:

$$\max_a \Pi(a;\theta) \qquad (6)$$

Because a FC opening does not always coincide perfectly with changes in the tax status of a state, we need to make an assumption about what Amazon believes entry into a state implies when it is making its profit maximizing decision. In our primary specification, we assume that Amazon believes that they will begin to collect sales tax immediately upon entry into a state. We have also estimated the model under alternative assumptions, which we discuss in Appendix C.

Given the maximization problem in 6, we know that the profit under the observed network rollout (a^o) must be greater than any deviation from this (a), or:

$$\Pi(a^o;\theta) \geq \Pi(a;\theta)\forall a \neq a^0$$

This allows us to formulate a set of moment inequalities that we use to estimate θ. [41]

Moment Inequalities

As in Holmes, we construct moment inequalities by comparing the discounted profits of the observed sequence of openings with counter-factual roll-out scenarios. We focus in particular on counter-factual roll-outs that *swap* the opening dates of two FCs (i.e., FC clusters). Each swap consists of moving the opening date of the earliest FC in the cluster. The swapped fixed costs are also assumed to be fixed costs (land and employees) associated with the earliest opening FC. The total number of perturbations is around 1,700 and after removing the swaps which open in the same year, it is 1,575.

By swapping two opening dates (e.g. $t < t'$), we ensure that the difference in the discounted stream of profit cancels out beyond period t'. This allows us to impose optimality conditions without observing the complete sequence of openings. Additionally, swapping the opening dates of two FCs implies that the unobserved sunk costs, S_t, will cancel out.

Consider two opening sequences: observed sequence a^0, and counter-factual sequence a. If profits were measured without error, profit maximization implies the following linear profit difference inequality:

$$y_a - x_a\theta \geq 0 \tag{7}$$

where $y_a = \sum_{t=2006}^{2018} \beta^{t-2006} \left[\sum_i \mu(R_{it}(a_t^0) - R_{it}(a_t)) - (F_{it}a_{it}^0 - F_{it}a_{it})\right]$ is the discounted profit differential net of shipping costs, $x_a = \sum_{t=2006}^{2018} \beta^{t-2006} \sum_i \left[d_{it}(a_t^0)R_{it}(a_t^0) - d_{it}(a_t)R_{it}(a_t)\right]$ is the discounted differences in revenue weighted shipping distance, and θ is the shipping cost per mile per dollar of revenue.

The cost savings associated with the optimal roll-out of fulfillment centers are identified by revealed preference. Intuitively, the magnitude of θ is such that the predicted revenue loss relative to counterfactual sequence a (i.e. $y_a < 0$) is offset by a reduction in shipping distance (i.e. $x_a < 0$). This gives a lower bound on the distance cost as it informs us of the value of θ such that the decrease in shipping distance exactly makes up for the loss in revenue. Similarly, θ is such that a counterfactual sequence associated with an increase in shipping distance (i.e. $x_a > 0$) is offset by a revenue gain from tax avoidance (i.e. $y_a > 0$). This provides an upper bound on the distance cost as it informs us of the value of θ such that the increase in shipping distance doesn't outweigh the

[41] Note that while we do not allow Amazon to re-optimize prices under a different network formation, the inequalities still hold because the profit under sub-optimal prices is weakly greater than the profit under optimal prices.

increase in revenue.

Although this tradeoff is clear when comparing y_a and x_a for alternative paths, we measure Amazon's profits and weighted shipping distance with error. In particular, we observe $\tilde{y}_a$ and $\tilde{x}_a$:

$$\begin{aligned}\tilde{y}_a &= y_a + \eta_a, \\ \tilde{x}_a &= x_a + \nu_a.\end{aligned}$$

These errors may originate from a variety of sources: (i) measurement error in the operating costs of FCs, (ii) estimation error in Amazon's revenue that would flow through to both y_a and x_a, (iii) omitted variables, and (iv) differences in tax implementation (i.e. Amazon's beliefs). We assume that both errors are independent of a vector of non-negative instruments z_a:

$$E(z_a\eta_a) = E(z_a\nu_a) = 0. \tag{8}$$

Applying these two moment restrictions to the "measured" inequality condition (analog from equation 7) gives:

$$E\left[z_a(\tilde{y}_a - \theta\tilde{x}_a)\right] = E[z_a(y_a - \theta x_a)] + \underbrace{E[z_a(\eta_a - \theta\nu_a)]}_{=0} \geq 0 \tag{9}$$

Thus, our estimator is based on a set of unconditional moment conditions consistent with the previous inequality:

$$E\left[z_a(\tilde{y}_a - \theta\tilde{x}_a)\right] \geq 0. \tag{10}$$

This leads to the following sample moment inequalities as in Pakes et al. (2015):

$$\frac{1}{M}\sum_a z_{a,k}(\tilde{y}_a - \theta^0\tilde{x}_a) = \tilde{m}_k(\theta^0) \geq 0, \quad \forall k = 1, \ldots, K. \tag{11}$$

where M denotes the number of inequalities.

Instruments

In order to construct the instruments we take a similar approach as Holmes (2011). That is, we form two sets of instruments. The first set include indicator variables that identify perturbations that capture the interaction between the effect of sales taxes and shipping distance. This corresponds to the interaction between cannibalization and economies of density in Holmes (2011). Here, we select inequalities based on how changes in shipping distance and/or taxes affect profit net of shipping cost and variable shipping cost (i.e. $\tilde{y}_a$ and $\tilde{x}_a$).

Specifically, we identify groups of counter-factual roll-outs that are useful in identifying the lower and upper bound of cost-savings. We select two types of "experiments" that reflect the revenue/shipping distance tradeoff. First, observing Amazon choosing to enter relatively early into a high-tax state when a low-tax (but further away) option was available identifies a lower bound of the shipping cost parameter. Therefore, to identify the lower bound, we look for counter-factual roll-outs that *increase the shipping distance* of the network to larger markets and result in higher revenues (Experiment 1). Similarly, observing Amazon choosing to enter relatively early into a low-tax state when a high-tax (but closer) option was available identifies an upper bound of the shipping cost saving parameter. Therefore, counter-factual roll-outs that *decrease the shipping distance* of the network to large markets identify an upper bound on the cost savings (Experiment 2). We group counter-factual rollouts that qualify for each experiment as follows.

1. **Experiment 1: Increase in shipping distance.** Let t_j be the chosen opening date of FC j and denote as S_2 the set of FCs opened at dates $t_{j'} > t_j$ in states that would lead to higher net revenue $\tilde{y}_{a(j,j')} < 0$. Define as $a(j, j')$ the counter-factual sequence that swaps j with $j' \in S_1$. Group all $j' \in S_1$ which have an increase in shipping distance:

    - Group 1: $j' \in S_1$ such that $\tilde{x}_{a(j,j')} < 0$

2. **Experiment 2: Decrease in shipping distance.** Let t_j be the chosen opening date of FC j and denote as S_2 the set of FCs opened at dates $t_{j'} > t_j$ in states that would lead to lower net revenue $\tilde{y}_{a(j,j')} > 0$. Define as $a(j, j')$ the counter-factual sequence that swaps j with $j' \in S_1$. Group all $j' \in S_2$ which have a decrease in shipping distance:

    - Group 1: $j' \in S_2$ such that $\tilde{x}_{a(j,j')} > 0$

These groups define a set of dummy variables that indicate whether a given counterfactual sequence that swaps observed opening date of FC j with that of FC j' qualifies for experiments 1 and 2, and if so, whether they have a positive or negative change in shipping distance. We denote this set of two dummy variables $\tilde{z}_a$. In practice, we could define more subgroups within each experiment (as in Holmes (2011)), but we have not found the results to be significantly different when taking this approach.

The second set of instruments is an interaction between these indicator variables and positive transformation of a continuous measures of the change in shipping distance[42]:

$$\tilde{x}_a^+ = \tilde{x}_a - \min_{a'} \tilde{x}_{a'}$$

A complication in our setting relative to Holmes (2011) is that both $\tilde{y}_a$ and $\tilde{x}_a$ are measured with error, meaning that we cannot use variables based on these as instruments. Specifically, Holmes (2011) uses an exogenous shifter of cannibalization (e.g., positive or negative changes in

[42] In alternative specifications presented in Appendix C we include the continuous variables without interacting them with the indicator variables. This allows us to use all the perturbations in estimation rather than limiting them to the perturbations that fall into experiment 1 or 2. Results are robust to this adjustment

store density) and changes in shipping distance, which is exogenous because shipping costs are assumed to be fixed (i.e., not a function of revenue), in order to define the indicators that serve as instruments. To get around this problem, we construct an alternative set of instruments by running a 'first-stage' regression of $\tilde{y}$ and $\tilde{x}$ on exogenous observables of perturbation a and then use the estimates of model to form the predicted values of $\tilde{y}$ and $\tilde{x}$, which we call $\hat{y}$ and $\hat{x}$. We then use $\hat{y}$ and $\hat{x}$ to form exogenous versions of the instruments, $\hat{z}_a$ and $\hat{x}_a^+$. Specifically, we run the regressions

$$\tilde{y}_a = \beta Z_a + \epsilon_a \tag{12}$$

$$\tilde{x}_a = \beta Z_a + \epsilon_a \tag{13}$$

where Z_a are exogenous shifters of $\tilde{y}_a$ and $\tilde{x}_a$ which include the overall changes in shipping distance, changes in sales tax rates, changes in fixed costs, and dummy variables indicating whether or not the perturbation features one or two FCs which are not the first to enter the state. We then use the estimates of β to form the predicted values of $\tilde{y}_a$ and $\tilde{x}_a$ which are then used to construct $\hat{z}_a$ and $\hat{x}_a^+$.

Estimation

We follow Andrews and Soares (2010) (AS hereafter) in order to construct our confidence sets for θ. For a given value of the shipping cost parameter θ_0 we calculate the sample moments $\bar{m}(\theta_0)$, which is a k $\times$ 1 vector, and the sample variance covariance matrix $\hat{\Sigma}(\theta_0)$ which is a k $\times$ k matrix. We then calculate the test statistic:

$$T(\theta_0) = S(n^{\frac{1}{2}}\bar{m}(\theta_0), \hat{\Sigma}(\theta_0))$$

where:

$$S(m, \Sigma) = \inf_{t=(t_1,0):t_1 \in R^p_{+,\infty}} (m-t)'\Sigma^{-1}(m-t) \tag{14}$$

We estimate $\hat{\Sigma}(\theta_0)$ using the method in Holmes (2011) in order to account for both first stage error and correlation across perturbations. See appendix B for details.

In order to calculate the critical value of the test statistic, we draw R bootstrap samples of perturbations. We sample N FCs from the population with replacement, where N is the observed number of FCs, to form a set which we denote N_r. We then construct a set of perturbations a_r from swapping the opening dates of the FCs contained in N_r. Note that this could result in duplicate swaps and/or swaps eliminated from the bootstrap sample. For each r, we form $S(\mu_r^{**}(\theta_0), \hat{\Omega}^{**}(\theta_0))$, where the starred variables are the boostrapped versions of the moments and the variance covariance matrix. See appendix B for details. We find the critical value, $\hat{c}(\theta_0, 1-\alpha)$, by the $1-\alpha$ sample quantile of $S(\mu_r^{**}(\theta_0), \hat{\Omega}^{**}(\theta_0))$ and we reject $H_0 : \theta = \theta_0$ if $T(\theta_0) > \hat{c}(\theta_0, 1-\alpha)$.

Using the bisection method, we find the maximum and the minimum value of θ such that

$T(\theta) \leq \hat{c}(\theta, 1-\alpha)$, which represent the upper bound and the lower bound of the confidence set, respectively. In practice, we set $R = 2000$.

Shipping Distance, Sales Tax, and Fixed Cost Tradeoff

We provide a description of the perturbations in order to demonstrate the trade-off Amazon faces between charging sales tax, shorter shipping distances and the fixed cost of operating FCs, which is the variation that identifies the shipping cost parameter. Figure 4 displays four scatter plots, where each point represents one of the 1,575 perturbations. The y-axis in the plots represents the difference in the population weighted average shipping distance between the observed network and the perturbed network. This varies from about 20 to -10 miles where most of the perturbations are between 5 and -5 miles. Each plot differs in its x-axis with Figure 4a using $\tilde{y}$, 4b using the difference in revenue times the revenue share, 4c using the difference in the number of households taxed and 4d using the difference in the fixed costs. The two lines displayed are the line of best fit for all the perturbations (red) and the line of bets fit for the perturbations which fall into experiments 1 and 2 (yellow).

There is a positive relationship between shipping distance and profit net of shipping cost ($\tilde{y}$), with the relationship being stronger for the 'selected' perturbations. This could be due to the fact that being closer to consumers implies that you are charging more of them sales tax and/or you are paying higher fixed costs. Figure 4b indicates that the former is true as there is positive relationship between revenue and shipping distance which is due to the fact there is a negative relationship between distance and the number of people taxed (see Figure 4c). The lines of best fit for these two plots are not significantly different, suggesting that the relationship between the sales tax implications of expansion and the shipping distance changes is similar for the selected perturbations and the ones which don't fall into experiments 1 and 2. Therefore, it is likely that the differences in the lines of best fit in Figure 4a are due to the fixed cost effects, something which is confirmed in Figure 4d.

Together the Figures in 4 suggest that Amazon faces a trade-off between charging sales tax and shipping distance, with fixed costs also playing an important role. Therefore, identification of the shipping cost parameter θ comes from both the tax/distance and the fixed cost/distance trade-off.

Results

We present the 95% confidence sets for θ under various definitions of the instruments in the first two panels of Table 10. We also show the number of perturbations which belong to each experiment.

We first focus on panel one, where we use the data to construct the instruments. The first row displays results when using just these indicator variables, while the second adds the interaction between these and $\tilde{x}^{+}$. We have explored including higher order functions of $\tilde{x}^{+}$, with little change in the estimated bounds. Note that the value of the parameter can be interpreted as the net cost of shipping one dollar of goods one mile. These estimates imply that the shipping cost is between 4.43E-05 and 1.04E-04, with the continuous instruments leading to a tighter interval.

Next, we turn to the second panel, where we display results when using a first stage regression to construct the instruments. Interestingly, this results in an upward shift of the confidence interval, suggesting that correlation in the error terms across perturbations leads to a downward bias in the estimates. The lower bound of the 95% confidence interval in these specifications is around 5.50E-05 and the upper bound is as high as 2.26E-04, with the addition of the continuous instruments again leading to a tighter interval. None of the confidence intervals include zero, suggesting that the cost savings from expansion are significantly different from 0.

The third panel of Table 10 estimates the model under the assumption that Amazon uses to FC with the lowest cost (as estimated from the shipping cost data) rather than the closest FC to serve a county. The confidence intervals become wider under this model, which is due to the fact that the correlation between shipping distance and the sales tax effects of expansion is not as strong, but the fact that they do not include zero provides further evidence that expansion has led to significant cost saving.

For the remainder of the analysis, we focus on the estimates that use both sets of instruments, with the instruments being constructed using the first stage regression (i.e., second row of panel 2). These results imply that it costs Amazon between \$0.58 and \$1.55 to ship \$100 of goods 100 miles, net of shipping revenue. Therefore, in 2006, it costs Amazon somewhere between \$1.74 to \$4.65 for the average \$100 shipment (300 miles) and between \$0.67 and \$1.78 in 2018 (115 miles), a significant amount of saves due to the expansion of the network.

Robustness

We compute three different sets of robustness estimates, which are all presented in more detail in Appendix C. First, we allow for an alternative assumption where there is a state specific lag between entry of a FC and the implementation of the tax law, and the lag is based on the observed lag. We base this assumption on the fact that many times the lag in the change in tax liability is due to Amazon negotiating (or challenging) the state government in terms of the laws. Therefore, the state level lag means that Amazon is correct in predicting the amount of time they will be able to avoid charging sales tax in a given state, either through negotiation or challenging the implementation of the law.

Second, we allow consumers to value the availability of same day shipping. Because same day shipping was implemented after our sample period, it may be the case that consumers value this service and that Amazon placed FCs taking this into consideration, which would be outside of our model. The opening of small FCs outside of major metropolitan areas is evidence that this may be the case. Because it would be difficult to directly estimate this demand side effect, we instead estimate the supply model when including an ad-hoc measure of the demand side effect of same day shipping while keeping the other demand parameters fixed at the estimated level.

Third, we allow for a more sophisticated model of Amazon logistics. The true shipping logistics model that Amazon employs is likely much more complicated than what we assume (i.e., coming from the closest FC). Therefore, in this model, we allow shipments to county i come from multiple FCs, where the amount coming from each FC is a function of shipping cost estimated using an outside data source and the capacity of the FC.

None of these adjustments change the results to a large degree, as we find a lower bound of the confidence interval as low as 1.06E-05 and an upper bound as high as 2.63E-04. This provides confidence that our model captures the important features of the data generating process and that our estimates are not systematically biased due to these assumptions.

Implications

We perform three exercises assuming that the shipping costs fall somewhere in between 5.75E-05 and 1.55E-04. First, we compute the total revenue, the total shipping costs and the total fixed costs over our sample period (2007-2018) under the observed network and under the 2006 network. This allows us to quantify the effect of the expansion over this time frame, holding all other features of the market, such as prices, fixed. The results are reported in the first four rows of Table 11. Predictably, Amazon loses revenue due to the tax effect and increases fixed operating costs as it expands its network, but it saves a significant amount in shipping costs. Specifically, the expansion led to a reduction in revenue of nearly \$9.6 billion, which is a loss \$2.2 billion in profit net of shipping cost, and an increase in fixed operating cost of \$3.3 billion. This is offset by between a \$5.0 and \$13.3 billion decrease shipping costs, which represents a 51% reduction.[43] Therefore, Amazon's expansion led to an increase of profit during this time period of between \$500 million and \$8.3 billion.

One weakness of this exercise is that it does not consider that Amazon might have problems fulfilling the number of orders under the 2006 network due to capacity constraints. Assuming that Amazon could adjust it's short-run capacity at a cost in order to fulfill orders, this implies that we are underestimating the total fulfillment costs under the 2006 network, meaning that our results are a lower bound on the total cost savings from expansion.

Another weakness of this exercise is that it does not consider the long term effects of expansion. While we cannot address this directly, we form a measure of the long term effect of expansion by computing a weighted average profit margin in 2018 under the observed network of FCs and then compare it to the 2018 margin under the network in 2006, where the weights are based on the

[43]The linearity of the shipping cost function results in the same percentage reduction for the upper and lower bounds

expenditures from each county. Specifically, we calculate:

$$\sum_i (\mu - \theta d_i(a))\omega_i$$

where a represents the network of FCs and ω_i is county i's share of aggregate 2018 revenue. The margin is 22.4 cents for every dollar sold under the lower bound of the confidence interval and 21.5 cents under the upper bound. Under the 2006 network, a^{2006}, the margin falls to between 21.3 cents and 18.4 cents per dollar, a reduction of between 1.2 and 3.1 cents for every dollar, or between 5 and 14%. Using our projections of Amazon's revenue in 2018, these estimates suggests that Amazon's 2018 profit alone is \$1.6 and \$4.5 billion dollars higher than it would be without expansion. One would think that this effect would only increase as Amazon continues to grow in revenue and FCs in the future. Again, these numbers could be thought of as a lower bound when considering that expansion also eases the costs of adjusting short term capacity.

Finally, we estimate how much the expansion has reduced the net average cost of shipping an order over the sample period. We calculate the average shipping cost by summing up the yearly fixed costs and the yearly shipping costs across all counties using our estimates of θ. We then divide these total costs by the number of orders in each year. One complication is that our model does not predict the number of orders. Instead, we rely on the raw ComScore data to create a yearly expenditures per order measure and then use the predicted expenditures from the model to calculate the expected number of orders per year. Figure 5 shows that the net average shipping cost falls from between \$2.67 and \$4.69 in 2006 to between \$1.09 and \$2.14 in 2013. Our model projects that this will continue to fall to approximately \$1 by 2018.

Put together, these three exercises demonstrate that Amazon's expansion has resulted in significant cost decreases as a result of the FC expansion.

Discussion

Using the results of both the demand and supply models, we provide some suggestive evidence as to the sources of the increase in concentration in online retail over the past decade. To do this, we decompose mode level preferences plotted in 3b into a price effect and a quality effect. Recall, that the fixed effect, $\hat{\bar{\xi}}_{jt}$, includes the effect prices, meaning that we can decompose the mode preferences into a price component, $\bar{p}_{jt}$, and a quality component, Ξ_{jt}:

$$\bar{\xi}_{jt} = \overbrace{\underbrace{(1-\hat{\sigma})log(p_{jt})}_{\bar{p}_{jt}} + \tilde{\xi}_{jt}}^{\hat{\bar{\xi}}_{jt}} \hspace{-2.5em} \underbrace{\phantom{\tilde{\xi}_{jt}} \hspace{0em} + \hat{\beta}_j \bar{Z}_t + \hat{\lambda}_j \bar{C}_t}_{\Xi_{jt}}$$

However, because our model does not predict anything about prices, we go to the raw data in order to form a 'price-index' for each mode and year.[44] Figure 6a plots these price indices measured absolutely (left axis) and in terns of their contribution to mode preferences (right axis). Amazon's prices have steadily declined from 2006-2013, while mode 3's have stayed relatively constant. Mode 2's followed Amazon's early on, only to level off from 2010 to 2013. In Figure 6b we graph the bounds on the average shipping cost over this time period. We can see that Amazon's costs decrease by approximately 55% (using the mid-point of the interval) and, over the same time period, their prices decreased by approximately 47%. While we cannot directly connect the cost decreases to the price decreases using our model, this provides suggestive evidence that expansion has helped Amazon achieve a competitive advantage.

Figure 6c displays the quality of each mode, which again shows a dramatic increase on Amazon compared to the other two modes. While there are many different components that could make up this quality, we examine one which we can measure with the raw data: product assortment. We calculate the HHI across product categories for each mode and year and plot it in Figure 6d. Amazon's assortment starts slightly higher than mode 2 in 2006, implying that they sell fewer product varieties than all the taxed sites put together. Mode 3 sells an even larger variety, which is not surprising considering that category is made up of the large collection of non-taxed online competitors. Over time, we see that Amazon has increased it's assortment a great deal and, by 2013, it had increased to the point where it was approximately equal to that of mode 3. Both mode 3 and mode 2 kept their assortment relatively constant over time. Again, while we cannot connect the patterns in these two graphs directly, this does suggest that at least some of the increase in Amazons share is due to an increase in assortment.

6 Conclusion

We examine sources for the increase in concentration of online retail over the last decade, with a focus on the economies of density associated with Amazon's large scale distribution network. Amazon benefits from their fulfillment centers being close to consumers for two potential reasons: first, the customer herself may value faster shipping, and second, it saves on delivery costs. At the same time, state laws dictate that Amazon must charge sales tax to consumers in most states where they have a physical presence. By raising the tax-inclusive price the consumer faces, this reduces consumers' willingness-to-pay for Amazon's services. Further, being close to significant population clusters raises the fixed cost of operating fulfillment centers. In order to analyze the connection between observed expansion and the increase in concentration, we estimate these different effects

[44] We run the following OLS regression:

$$p_{ocjt} = \xi_{jt} + \lambda_c + \epsilon_{ocjt}$$

where the dependent variable is the tax exclusive price on purchase occasion o, in product category c, on mode j, an in year t. We use the estimated value of ξ_{jt} as the price index of mode j in year t.

of expansion.

We find that Amazon indeed faces this tradeoff: consumers dislike paying taxes, meaning there be must be gains of network expansion due to faster shipping and/or reduced delivery costs. Our demand estimates indicate that consumer demand does not respond to our measures of shipping times, most likely due to the fact that expansion did not actually result in faster shipping speeds for the vast majority of consumers apart from the possibility of one-day shipping. Therefore, we find that the network expansion from 2006-2018 resulted in a loss in revenue of around \$9.6 billion dollars for Amazon, but at the same time reduced the average shipping distance from FC to consumer by around 180 miles by the end of 2018. We use a moment inequalities approach, together with the assumption that Amazon's network expansion path is optimal, to infer shipping cost savings from the observed fulfillment center network relative to alternative configurations. Results suggest that Amazon saves between \$0.58 and \$1.55 per 100 miles for every \$100 dollars of goods shipped. Therefore, the expansion of the network has resulted in between \$5 and \$13.3 billion in savings on shipping costs and an increase in profit margins of up to 14%.

While our model doesn't allow for a direct connection between these costs and concentration, we show that the decrease in costs is mirrored by a decrease in prices relative to other online retailers thus providing suggestive evidence of the role of economies of density. Further, we show that improvements in the quality of Amazon's platform is another driver of the increase in concentration, with an increase in product assortment likely being a source for this improvement.

References

Agarwal, S., S. Chomsisengphet, T. Ho, and W. Qian (2013). Cross border shopping: Do consumers respond to taxes or prices? Working Paper.

Agarwal, S., N. Marwell, and L. McGranahan (2013). Consumption responses to temporary tax incentives: Evidence from state sales holidays. Working Paper.

Agatz, N. A., M. Fleischmann, and J. A. Van Nunen (2008). E-fulfillment and multi-channel distribution–a review. *European Journal of Operational Research 187*(2), 339–356.

Alm, J. and M. I. Melnik (2005). Sales taxes and the decision to purchase online. *Public Finance Review 33*(2), 184–212.

Anderson, E. T., N. M. Fong, D. I. Simester, and C. E. Tucker (2010). How sales taxes affect customer and firm behavior: The role of search on the internet. *Journal of Marketing Research 47*(2), 229–239.

Andrews, D. W. and G. Soares (2010). Inference for parameters defined by moment inequalities using generalized moment selection. *Econometrica 78*(1), 119–157.

Asplund, M., R. Friberg, and F. Wilander (2007). Demand and distance: evidence on cross-border shopping. *Journal of Public Economics 91*(1), 141–157.

Ballard, C. L. and J. Lee (2007). Internet purchases, cross-border shopping, and sales taxes. *National Tax Journal 60*(4), 711–725.

Baugh, B., I. Ben-David, and H. Park (2014). The "Amazon Tax": Empirical Evidence from Amazon and Main Street Retailers. Technical report, National Bureau of Economic Research.

Bruce, D., W. F. Fox, and L. Luna (2009). State and local government sales tax revenue losses from electronic commerce. *State Tax Notes 52*(7), 537–558.

Brynjolfsson, E., Y. J. Hu, and M. D. Smith (2003). Consumer surplus in the digital economy: Estimating the value of increased product variety at online booksellers. *Managment Science 49*(11), 1580–1596.

Chetty, R., A. Looney, and K. Kroft (2009). Salience and taxation: Theory and evidence. *The American Economic Review 99*(4), 1145–1177.

De los Santos, B., A. Hortaçsu, and M. R. Wildenbeest (2012). Testing models of consumer search using data on web browsing and purchasing behavior. *The American Economic Review 102*(6), 2955–2980.

Einav, L., J. Levin, and N. Sundaresan (2014). Sales taxes and internet commerce. *The American Economic Review 104*(1), 1–26.

Ellison, G. and S. F. Ellison (2009). Tax sensitivity and home state preferences in internet purchasing. *American Economic Journal: Economic Policy 1*(2), 53–71.

Goolsbee, A., M. F. Lovenheim, and J. Slemrod (2010). Playing with fire: Cigarettes, taxes, and competition from the internet. *American Economic Journal: Economic Policy 2*(1), 131–154.

Holmes, T. J. (2011). The diffusion of Wal-Mart and economies of density. *Econometrica 79*, 253–302.

Pakes, A., J. Porter, K. Ho, and J. Ishii (2015). Moment inequalities and their application. *Econometrica 83*(1), 315–334.

Quan, T. and K. Williams (2016). Product Variety, Across-Market Demand Heterogeneity, and the Value of Online Retail. Working Paper.

Scanlan, M. A. (2007). Tax Sensitivity in Electronic Commerce. *Fiscal Studies 28*(4), 417–436.

Smith, M. D. and E. Brynjolfsson (2001). Consumer decision-making at an internet shopbot: Brand still matters. *Journal of Industrial Economics 49*(4), 541–558.

Zheng, F. (2014). Spatial Competition and Preemptive Entry in the Discount Retail Industry. Working Paper.

Table 1: Expansion of fulfillment center network

Year	FCs	States with FC	States Taxed	Counties Taxed	Ave Distance (miles)	Avg FC Size (1,000 sqft)	Avg FC Employees
2006	8	6	4	317	297	544	504
2007	9	7	4	317	292	522	484
2008	12	10	5	379	236	487	452
2009	17	10	5	379	227	570	529
2010	17	10	5	379	243	570	529
2011	24	10	5	379	236	657	610
2012	32	12	8	752	223	718	666
2013	41	14	16	1,189	207	765	709
2014	48	14	23	1,644	175	707	656
2015	54	16	26	1,937	152	665	616
2016	90	27	27	1,983	123	562	521
2017	101	28	28	2,047	116	552	511
2018	104	28	28	2,047	116	558	517

Notes: The number of states where Amazon purchases are subject to sales tax exceeds the number of states with a FC due to states negotiating for sales taxes being collected immediately after agreeing to let Amazon build a FC in the state, even if there is a delay between the time of agreement and the actual opening of the warehouse. Average distance is the weighted average distance from a county centroid to the nearest Amazon FC, where the weights are based on county population.

Table 2: Household Online Purchasing

Year	Households (1,000)	Counties (%)	Expenditures	Transactions	Offline Only (%)	Adjusted Expenditures	Adjusted Transactions	Offline Only (%)
2006	87.1	92	$248	2.5	49.8%	$250	2.5	55.5%
2007	90.0	92	$260	2.6	51.2%	$248	2.4	60.8%
2008	57.0	87	$201	2.1	59.3%	$274	2.8	-
2009	55.9	85	$148	1.5	66.7%	$282	2.9	-
2010	54.1	84	$129	1.4	67.9%	$288	3.1	32.1%
2011	63.4	87	$134	1.4	69.3%	$342	3.5	23.0%
2012	55.2	84	$155	1.8	63.6%	$327	3.8	23.9%
2013	46.5	78	$164	2.2	60.1%	$374	5.0	15.5%

Notes: Displayed are the number of households in the comScore data and the percentage of counties represented. Online expenditure and transactions and the share of households with zero online expenditures come from comScore, while the adjusted variables are calculated following the procedure in Appendix A. The percent of offline shoppers only denotes the share of respondents who answered no to the question whether they had shopped online in the previous three months in the Forrester Technographics Survey.

Table 3: Transaction Level Demand Estimates

Variable name	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Tax Elasticity	-0.138	-0.124	-0.139	-0.134	-0.144	-0.145	-0.144
	(0.073)	(0.074)	(0.074)	(0.073)	(0.075)	(0.074)	(0.075)
Local Express Delivery		-0.016					
		(0.010)					
Log Distance			0.000				
			(0.002)				
1 or 2 Day Priority				0.019			
				(0.011)			
1, 2, or 3 Day Package					-0.004		
					(0.004)		
1 Day First Class						-0.004	
						(0.004)	
1, 2, or 3 Day Standard							-0.004
							(0.004)
Obs	2,291,291	2,291,291	2,291,291	2,291,291	2,291,291	2,291,291	2,291,291
R-Sq	0.354	0.354	0.354	0.354	0.354	0.354	0.354

** 1% * 5%. Standard errors are clustered at the county/year level. Product category and year fixed effects included in all models. Shipping times are grouped into "long" and "short" with "long" being the excluded category.

Table 4: Diff-in-Diff Demand Estimates

Variable name	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Taxed Dummy	-0.095	-0.095	-0.097	-0.094	-0.106*	-0.103*	-0.106*
	(0.051)	(0.051)	(0.051)	(0.051)	(0.052)	(0.052)	(0.052)
Local Express Delivery		-0.008					
		(0.159)					
Log Distance			0.012				
			(0.029)				
1 or 2 Day Priority				0.010			
				(0.090)			
1, 2, or 3 Day Package					-0.041		
					(0.035)		
1 Day First Class						-0.043	
						(0.044)	
1, 2, or 3 Day Standard							-0.041
							(0.035)
Obs	12,830	12,830	12,830	12,830	12,830	12,830	12,830
R-Sq	0.362	0.362	0.362	0.362	0.362	0.362	0.362

** 1% * 5%.Notes: Regressions include year and county fixed effects. Shipping times are grouped into "long" and "short" with "long" being the excluded category.

Table 5: Taxed and Non-Taxed Competitors

Sales Rank	Taxed	Non-Taxed
1	walmart.com	dell.com
2	jcpenney.com	qvc.com
3	staples.com	yahoo.net
4	victoriassecret.com	hsn.com
5	officedepot.com	yahoo.com
6	bestbuy.com	quillcorp.com
7	apple.com	overstock.com
8	target.com	ebay.com
9	sears.com	orientaltrading.com
10	costco.com	zappos.com
Total (%)	192 (34)	375 (66)

Notes: Table displays top 10 domains that we define as taxed and non-taxed.

Table 6: CES Demand Estimates

Variable name	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Tax Elasticity	-1.430**	-1.454**	-1.471**	-1.440**	-1.400**	-1.429**	-1.400**
	(0.486)	(0.486)	(0.490)	(0.486)	(0.486)	(0.486)	(0.486)
Local Express Delivery		0.186					
		(0.149)					
Log Distance			-0.011				
			(0.016)				
1 or 2 Day Priority				-0.058			
				(0.058)			
1, 2, or 3 Day Package					0.040		
					(0.031)		
1 Day First Class						0.002	
						(0.034)	
1, 2, or 3 Day Standard							0.040
							(0.031)
Obs	42,390	42,390	42,390	42,390	42,390	42,390	42,390
R-Sq	0.236	0.236	0.236	0.236	0.236	0.236	0.236

** 1% * 5%. Notes: Presented are the estimates of tax sensitivity and the effect of shipping speeds. Regressions include mode/year dummies along with mode level effects of local demographics and competition. Shipping times are grouped into "long" and "short" with long being the excluded category.

Table 7: Demand Estimates (Mode Fixed Effects)

Year	Amazon	Other Taxed Online	Non-taxed Online
2006	-1.777**	-0.597	
	(0.343)	(0.322)	
2007	-1.421**	-0.303	0.191*
	(0.352)	(0.332)	(0.084)
2008	-0.850*	0.088	0.461**
	(0.354)	(0.334)	(0.091)
2009	-0.630	0.123	0.464**
	(0.355)	(0.334)	(0.095)
2010	-0.578	-0.020	0.326**
	(0.356)	(0.335)	(0.098)
2011	-0.071	0.338	0.496**
	(0.354)	(0.334)	(0.095)
2012	0.220	0.536	0.684**
	(0.352)	(0.333)	(0.094)
2013	0.663	0.711*	0.934**
	(0.354)	(0.335)	(0.102)

** 1% * 5%. Notes: Presented are the estimated mode/year effects from specification (1) in Table 6.

Table 8: Demand Estimates (Demographics)

Varaible	Amazon	Mode 2	Mode 3
Total Offline Competitors	-0.002	-0.002	-0.003
	(0.016)	(0.016)	(0.016)
Large Offline Competitors	-0.010	-0.016**	-0.010
	(0.006)	(0.006)	(0.006)
Income	0.655**	0.638**	0.629**
	(0.020)	(0.018)	(0.018)
% Pop Black	-0.309	0.328	
	(0.271)	(0.241)	
% Pop White	-0.130	0.442	
	(0.258)	(0.253)	
% Pop Asian	1.660*	-0.386	
	(0.821)	(0.800)	

** 1% * 5%. Notes: Presented are the estimated demographic effects from specification (1) in Table 6. The offline competition variable is measured in hundreds of establishments.

Table 9: Predicted Revenue ($Billion)

Year	Amazon	Mode 2	Mode 3	Total
2006	5.61	21.21	25.98	52.79
2007	7.77	27.76	30.65	66.18
2008	9.78	27.81	29.01	66.60
2009	12.28	29.06	29.37	70.71
2010	17.88	34.91	35.38	88.17
2011	25.27	42.56	35.73	103.57
2012	32.46	50.22	41.82	124.50
2013	40.79	51.23	43.22	135.24
2014	53.59	61.70	49.14	164.44
2015	65.85	66.57	52.33	184.74
2016	80.04	71.30	55.43	206.78
2017	96.09	75.90	58.44	230.43
2018	114.11	80.33	61.34	255.78
Total	561.53	640.55	547.82	1749.91

Notes: Presented in the first panel are the revenue predictions based on the model estimates. Future revenues are calculated based on projections of the mode/year and region/year effects.

Table 10: Cost Estimates

Instruments	95% Confidence Set	# Perturbations (Group 1)	# Perturbations (Group 2)
	Base: No First Stage		
$\tilde{z}_a$	[3.43E-05,1.04E-04]	332	496
$\tilde{z}_a, \tilde{z}_a \times \tilde{x}_a^+$	[3.60E-05,9.29E-05]	332	496
	Base: First Stage		
$\hat{z}_a$	[5.45E-05,2.26E-04]	263	564
$\hat{z}_a, \hat{z}_a \times \hat{x}_a^+$	[5.75E-05,1.55E-04]	263	564
	Lowest Cost FC: First Stage		
$\hat{z}_a$	[3.90E-05,1.20E-04]	180	557
$\hat{z}_a, \hat{z}_a \times \hat{x}_a^+$	[3.50E-05,4.56E-04]	180	557

Notes: Confidence sets computed using the methods of Andrews and Soares (2010). The instruments indicated by $\hat{z}$ are determined using a first-stage regression. The last two columns indicate the number of perturbations that belong to each experiment. The third panel assumes that Amazon uses the lowest cost FC, rather than the closest, to serve the customers in a county.

Table 11: Effects of Expansion ($Billion)

	Observed Network	2006 Network	Change	% Change
μ*Revenue	129.15	131.36	-2.21	-1.71%
Labor Cost	4.30	1.41	2.90	67.34%
Land Cost	0.57	0.14	0.43	74.94%
Shipping Cost	[4.66 , 12.55]	[9.62 , 25.91]	[-4.96 , -13.36]	-51.56%
2018 Margin	[0.224 , 0.215]	[0.213 , 0.184]	[0.012 , 0.031]	[5.13% , 14.43%]

Notes: All figures are measured in billions of dollars except where noted. Table compares outcomes under observed FC rollout to outcomes holding the FC network fixed as in 2006. For this counterfactual FC network, revenue is predicted assuming that tax status of a county does not change from 2006 onwards and that shipping distance remains fixed. The cost savings results are the 95% confidence interval using the estimates from Table 10.

Figure 1: Concentration in Online Retail, 2006-2013

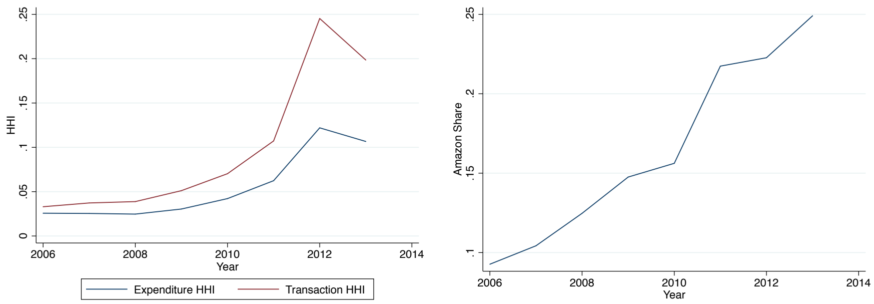

(a) HHI (b) Amazon Share

Notes: Figures are based on online sales from product categories that Amazon sells.

Figure 2: Online Competition

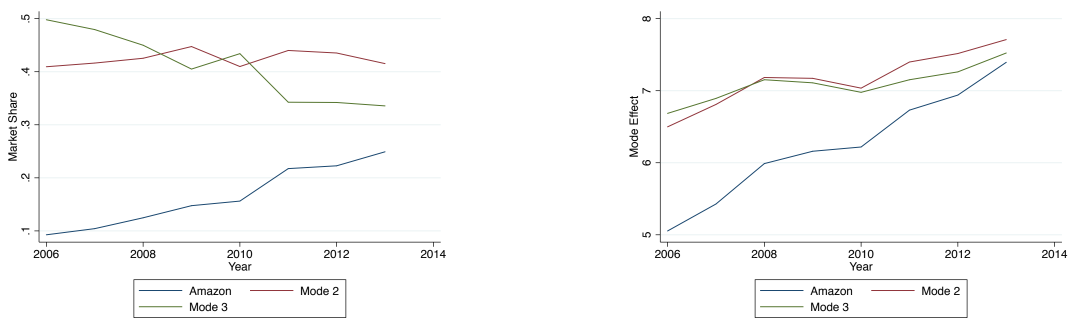

(a) Market Share (b) Mode Effects

Notes: Market shares are calculated using adjusted expenditures and the mode effects come from the estimates of the demand model.

Figure 3: Perturbations

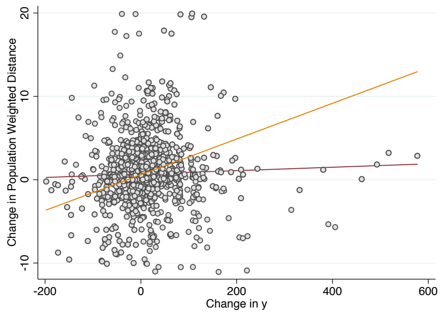

(a) $\tilde{y}$

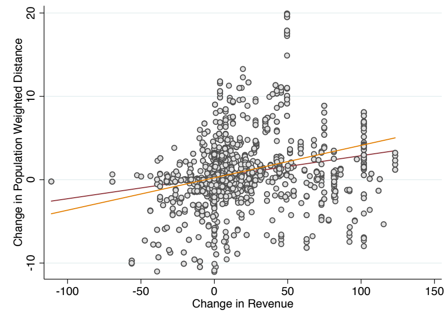

(b) Change in $\mu * R$

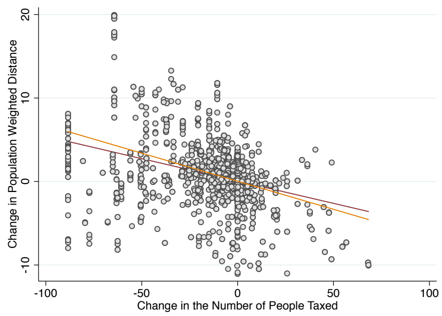

(c) Change in Population Taxed

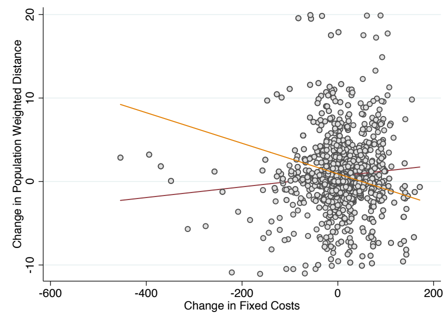

(d) Change in Fixed Costs

Notes: Each point represents one of the 1,575 perturbations. The y-axis in all four pictures is the difference in the weighted average shipping distance between the observed network and the perturbation, where the weights are based on the the number of households the FC serves. The red line is the line of best fit for all the perturbations and the yellow line is the line of best fit for only the perturbations which fall into experiment 1 or 2.

Figure 4: Average Shipping Cost

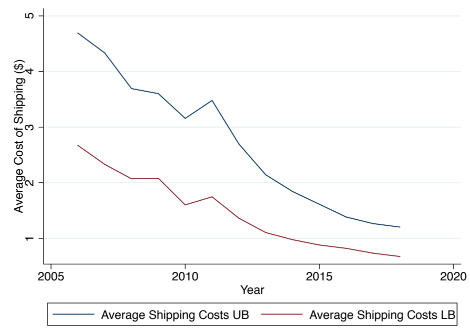

Notes: The figure displays the 95% confidence interval on the average cost of shipping an order using the estimates from the fourth row of 10. Total orders are inferred from a combination of the raw data and predictions of the model.

Figure 5: Sources of Concentration

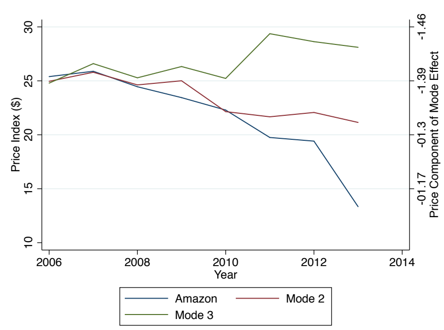

(a) Prices

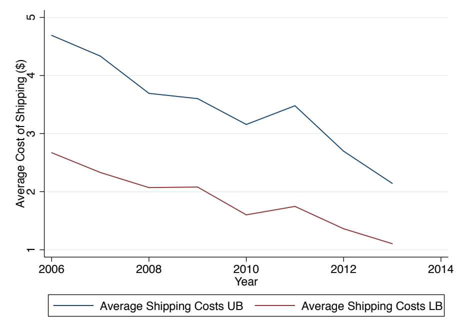

(b) Average Cost

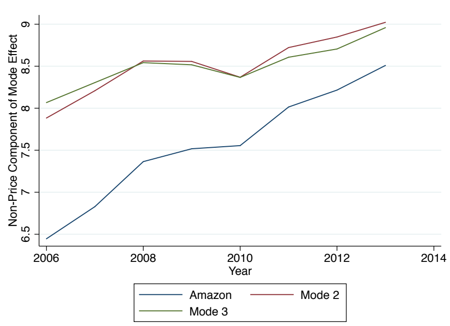

(c) Quality

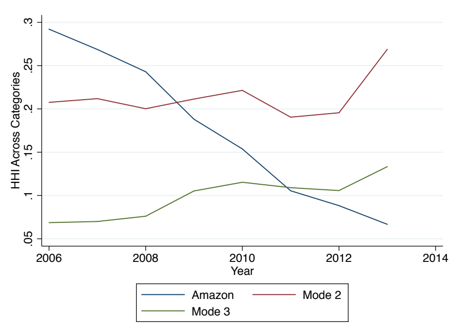

(d) Assortment

Notes: Price indices are calculated using the raw comScore data netting out any category effects. Quality is the calculated by removing the price component from the estimated mode effect. Assortment is measured by the HHI across product categories within each mode.

A Data Appendix

Forrester Data

We purchased survey data from Forrester Research, Inc which provides information on the extent of online shopping.[45] The survey, called the "North American Technographics Online Benchmark Survey", was conducted from 2006-2007 by mail and 2010-2014 online and surveyed between 30,000 and 60,000 households in the United States and Canada. The exact content of the questions on the survey varies by year, but generally the most pertinent question for us is the one which asks the user "Have you bought anything online in the past three months?". Additionally, the survey asks for information about the age, income, race, and zip code of the respondent. The documentation of the survey is available from the authors upon request.

Calculating Expenditures

In what follows, we provide a description of the procedure to calculate the expenditures for the representative consumer in county i and year t on the three different modes of shopping.

We first use the Forrester data to estimate the extent of online shopping for a given demographic group that we then use to account for the extensive online shopping margin, which is measured poorly in the ComScore data. We run the following linear probability regression:

$$Pr(D_{\iota t} = 1) = \beta_0 + \beta_1 Z_\iota + \beta_2 I_{iota} + \gamma_t + \epsilon_{\iota t}$$

where Z includes set of dummy variables indicating the characteristics of respondent ι (i.e., their race, age and census region) and a continuous measure of their income, and γ_t is a year dummy variable. The dependent variable, $D_{\iota t}$, represent whether or not the consumer indicated that they had purchased anything online in the past three months. We then use the estimates of the model to predict the probability that a consumer who belongs to group Z purchased something online in year t. That is, we form $\hat{p}_{zt}$ for each consumer group and year. For the years 2008 and 2009, we use linear interpolations based on the predictions in 2007 and 2010.

Using these predicted probabilities, we calculate the Forrester adjusted expected total online expenditures for a household who belongs to group z in county i and year t:

$$e_{zit} = \frac{1}{N_{it}(z)} \sum_{h \in H_{it}(z)} \hat{p}_{zt} \tilde{e}_{ht}$$

Here, $\tilde{e}_{ht}$ is the observed expenditure of household h, $H_{it}(z)$ is the set of households in group z in county i, and $N(z)_{it}$ is the size of this set. Importantly, we drop any consumer where expenditures are equal to 0, because this is accounted for through the probability of online shopping.

[45] `https://www.forrester.com/home/`

Next, we calculate group level shares for each shopping mode in each county using the unadjusted ComScore data:

$$s_{zitj} = \frac{\bar{e}_{zitj}}{\sum_j \bar{e}_{zitj}}$$

where j indicates the shopping mode and $\bar{e}_{zitj}$ is the average expenditures on mode j in county i for group z:

$$\bar{e}_{zitj} = \frac{1}{N_{it}(z)} \sum_{h \in H_{it}(z)} \tilde{e}_{hitj}$$

Using census weights for demographic group z, denoted w_{zi}, and the adjusted ComScore data, we create the total online expenditures for a representative consumer in a given county:

$$e_{it} = \sum_z w_{zi} e_{zit}$$

and the representative shares:

$$s_{itj} = \sum_z w_{zi} s_{zitj}$$

Finally, to get the expenditures for the representative consumer across each mode, we combine the previous two calculations:

$$e_{itj} = s_{itj} \tilde{e}_{it}$$

B Estimation

Estimation of Variance-Covariance Matrix

In our problem, we cannot simply calculate $\hat{\Sigma}(\theta_0)$ using equation 3.2 in AS because (1) there is first stage error in our revenue equation and (2) there exists correlation across a because the same FCs are used in many perturbations. We take the approach described in Section 8.1 of Holmes (2011) in order to account for these two issues .

Specifically, we draw a subsample of FCs of size $b << N$, where N is the total number of FCs. We keep the perturbations in which both the swapped FCs are in the subsample, thereby forming the 'deviation subsample' indexed by s. We then take a draw of the tax parameter $\hat{\sigma}_s$ from a normal distribution centered at $\hat{\sigma}$ with variance $\frac{N}{b}\hat{\Sigma}^{\sigma}$, where $\hat{\Sigma}^{\sigma}$ is the estimated variance from the first stage revenue model. Note that in order to account for first stage error, we only draw the tax parameter and assume that all other features of the demand model approximately cancel out when calculating the perturbation. This assumption is made for simplification because for each draw of the parameters, we must compute all the perturbations. There are over 90 parameters of the demand model, meaning we would have to draw a large number of vectors of parameters in order to have variation across the full joint distribution. Instead, we take 100 draws of the tax parameters and calculate the set of perturbations for 100 times. While the non-linearity of the demand model imply that the other components of demand won't exactly cancel out, we believe that it is a good approximation and that the results will not be significantly affected.

We calculate $\tilde{y}_s$ and $\tilde{x}_s$ with the deviation subsample and the new tax parameter σ_s to form:

$$\bar{m}_{s,k}(\theta_0) = \frac{1}{M}\sum_{a} z_{a,k}[\tilde{y}_s - \theta_0\tilde{x}_s]$$

resulting in the k $\times$ 1 vector $\bar{m}_s(\theta_0)$). The variance-covariance matrix is then estimated from the S different deviation subsamples:

$$\hat{\Sigma}(\theta_0) = \frac{b}{N}\text{varcov}(\bar{m}(\theta_0))$$

where $\bar{m}(\theta_0)$ is the S $\times$ k matrix of moments.

In practice, we set $b = \frac{N}{3}$ and $S = 200$ and use the diagonal matrix equivalent of $\hat{\Sigma}(\theta_0)$ as our weight matrix in equation 14. The reason for the latter is that we ran into a number of instances where the full matrix was not invertible.

Bootstrap Sample Statistic

We follow the procedure is described in Section 4.2 of Andrews and Soares (2010). From the bootstrap sample of R perturbations, we calculate $\{\mu^*_r(\theta_0), \hat{\Omega}^*(\theta_0)\}$ where:

$$\mu^*(\theta_0) = M_r^{\frac{1}{2}}(\hat{D}^*(\theta_0))^{-\frac{1}{2}}(\bar{m}^*(\theta_0) - \bar{m}(\theta_0))$$

$$\hat{\Omega}^*(\theta_0) = (\hat{D}^*(\theta_0))^{-\frac{1}{2}}\hat{\Sigma}^*(\theta_0)(\hat{D}^*(\theta_0))^{-\frac{1}{2}}$$

The elements of $\bar{m}^*(\theta_0)$ are given by:

$$\bar{m}^*_k(\theta_0) = \frac{1}{M_r}\sum_{a_r} z_{a_r}[\tilde{y}_{a_r} - \theta_0\tilde{x}_{a_r}]$$

and:

$$\hat{D}^*(\theta_0) = Diag(\hat{\Sigma}^*(\theta_0))$$

$Diag(X)$ represents a vector containing the diagonal elements of X. Here, $\hat{\Sigma}^*(\theta_0)$ is computed using the same subsampling method described above, but replacing the the 'population' of FCs from which we draw the deviation subsample with N_r. We then eliminate the elements of $\{\mu^*_r(\theta_0), \hat{\Omega}^*(\theta_0)\}$ where:

$$M^{\frac{1}{2}}\frac{\bar{m}_k(\theta_0)}{\hat{\nu}_k(\theta_0)} > \kappa = (ln(M))^{\frac{1}{2}}$$

where $\bar{m}_k(\theta_0)$ is the kth element of $\bar{m}(\theta_0)$ and $\hat{\nu}_k(\theta_0)$ is the corresponding standard deviation taken from $\hat{\Sigma}(\theta_0)$. This results in a weakly smaller set of moments and their variance-covariance matrix:

$$\{\mu_r^{**}(\theta_0), \hat{\Omega}^{**}(\theta_0)\}$$

Using these elements, we calculate $S(\mu_r^{**}(\theta_0), \hat{\Omega}^{**}(\theta_0))$ via equation 14.

Calculating μ

Amazon reports the total amount of revenue in the "Media" and "Electronics and Other General Merchandise" categories in North America, which is roughly equivalent to the revenue we predict from our model. Revenue from Canada and Mexico is not included in the ComScore data. Therefore, these sales are accounted for through the 'multiplier' that we use to match the predictions of the model in the financial reports. See Section 4 for a discussion of these multipliers. They also report the "Cost of Goods Sold" for all of their sales. We compute the cost of goods sold for North America by multiplying the total cost by the ratio of sales from North America to total sales. This provides us with a value for the "gross margin". However, as Amazon states, the reported cost of goods sold includes both outbound shipping costs and inbound shipping costs (through wholesale prices). Recall that we assume that inbound shipping costs from suppliers to Amazon do not vary with the network of FCs. As we are estimating the outbound shipping costs, we exclude them from the gross margin by adding in the reported "Net Shipping Cost as a Percentage of Revenue" to the gross margin. This value grows over time, but we suspect this is due to an increase in Amazon's non-shipping related activities (i.e. Amazon Web Services and digital goods). We therefore set $\mu = 0.23$.

C Robustness

Demand

In what follows, we provide robustness to the demand estimates reported in the body of the paper. Tables 12 and 13 provide results for a number of different variants of the demand model and results when adjusting the sample used in estimation, while Table 14 displays results under alternative measures of the shipping distance.

In Table 12, specification (1) uses weighted OLS, where the weights are analytical weights based on the number of observations used to calculate the average expenditures in a county, specification (2) incorporates the zeros by changing any mode-level expenditures that are equal to \$0 to \$1, specification (3) uses data from 2008 and beyond because the number of households in the comScore sample shifts after 2007, specification (5) does not use the Forrester correction to adjust expenditures, and specification (6) does not use population weights to create the expenditures.

Specification (4) estimates the tax sensitive separately for each mode, as a way of accounting for Amazon Marketplace. That is, it allows for an Amazon specific tax effect that may be lower than the effect for offline retailers and mode 3 because not all customers on Amazon are charged sales tax.

In Table 13, the first two specifications include a state-year level fixed effect (1) or a county-year level fixed effect (2) rather than a region-year level effect in order to account for the possibility that the unobservable factors which may be correlated with the tax variable are more local compared to our main specification. Specification (3) includes dummy variables indicating whether or not a FC entered the state in which the county is located one, two, three, or more than three years ago. The is intended to capture any sort of delayed marketing push or push to get users to adopt Amazon Prime which is associated with the entry of a FC. Finally, specification (4) includes an interaction between these dummy variables and the tax variable in order capture any possible learning which may happen over time. That is, it may take consumers time in order to learn that they are actually being charged higher prices than the were before, meaning the response to FC entry may be delayed. While we don't display the results here, we have also explored a learning effect with regards to shipping speeds, but again fail to find any evidence that our measures of shipping speeds affect consumer decisions.

In Table 14, we adjust the measure of shipping speed by allowing the consumer from county i to receive shipments from different FCs, where the amount of good coming from each FC is a function of the estimated shipping cost from that FC and the capacity. See the Cost section of this Apopendix for details on how we derive the distribution of goods coming from different FCs.

Overall, the results of these robustness checks are consistent with the base results with the estimated tax sensitivity being between -1.2 and -1.9 and the convenience effect not being significant.

Table 12: Robustness (Demand 1)

Variable name	(1)	(2)	(3)	(4)	(5)	(6)
Tax Elasticity	-1.307**	-1.687**	-1.867**		-1.203	-1.199**
	(0.286)	(0.585)	(0.575)		(0.625)	(0.401)
Tax Elasticity (Amazon)				-1.166*		
				(0.515)		
Tax Elasticity (Mode 3)				-2.900**		
				(0.847)		
Obs	42,399	52,617	29,053	42,399	43,811	42,400
R-Sq	0.315	0.162	0.135	0.185	0.137	0.199
Regression	A-Weights	Zeros	2008-2013	Individual Tax Effect	No Forr Adjustment	No Pop Weights

** 1% * 5%. Notes: Presented are the robustness estimates excluding any shipping speed effect. The results when including these effects are similar to that of the base regressions.

Table 13: Robustness (Demand 2)

Variable name	(1)	(2)	(3)	(4)
Tax Elasticity	-1.325**	-1.215**	-1.401**	-1.644**
	(0.498)	(0.443)	(0.481)	(0.525)
Entry Year 0			-0.034	
			(0.050)	
Entry Year -1			-0.015	
			(0.061)	
Entry Year -2			0.093	
			(0.072)	
Entry Year -3			0.046	
			(0.073)	
Entry Year < -3			-0.078	
			(0.058)	
Tax*(Entry Year 0)				-0.620
				(0.759)
Tax*(Entry Year -1)				-0.045
				(0.845)
Tax*(Entry Year -2)				-0.719
				(1.061)
Tax*(Entry Year -3)				-1.104
				(1.082)
Tax*(Entry Year < -3)				0.551
				(0.406)
Obs	42,399	42,399	42,399	42,399
R-Sq	0.195	0.240	0.186	0.186
Regression	Year/State Effects	Year/County Effects	Years Since Entry	Years Since*Tax

** 1% * 5%. Notes: Presented are the robustness estimates excluding any shipping speed effect. Specifications (1) and (2) include different delineations of the fixed effects, while (3) includes dummy variables indicating the time since the FC entered a consumer's state and (4) includes interactions between these dummies and the tax rate charged.

Table 14: Robustness (Demand 3)

Variable name	(1) {10, 0}	(2) {10, 20}	(3) {20, 0}	(4) {20, 20}	(5) {200, 0}	(6) {200, 20}	(7) Low Cost
Tax Elasticity	-1.469**	-1.497**	-1.468**	-1.497**	-1.508**	-1.531**	-1.474**
	(0.481)	(0.480)	(0.480)	(0.480)	(0.494)	(0.487)	(0.480)
Log Distance	-0.031	0.025	-0.028	0.016	-0.015	-0.022	-0.020
	(0.040)	(0.025)	(0.030)	(0.026)	(0.025)	(0.023)	(0.024)
Obs	42,399	42,399	42,399	42,399	40,755	41,695	42,399
R-Sq	0.186	0.186	0.186	0.186	0.185	0.186	0.186

** 1% * 5%. Notes: Presented are the robustness estimates including different measures of the shipping speed, where the distance is calculated based on the logistics models discussed below in C. The numbers below the specification number indicate the values of β_1 and β_2.

Cost: State Level Lag

We allow for the assumption about when the tax laws will change upon entry to be a function of the observed lag between entry and the change in the law as described in Section 5. For example, if we swap a FC that opened in Pennsylvania in 2008 and resulted in sales tax liabilities starting in 2011 with one that opened in Nevada in 1999 but resulted in sales tax liabilities only in 2014, then we would assume that the FC in Pennsylvania opens in 1999 and sales tax is collected beginning in 2002 (i.e., there is a 3 year lag as there was with the actual FC that opened in Pennsylvania in 2008).

The results in Table 15 are mostly in line with the no-lag assumption, with slightly wider confidence intervals for the base model and tighter intervals for the low-cost model. Importantly, non of the confidence sets include zero.

We have also estimated the model under a opening date specific lag. In the example above, this assumption would imply that the lag would "follow" the swapped FCs. That is, taxes would begin being collected in Nevada in 2011 and in Pennsylvania in 2014. The results under this specification do not vary significantly from the estimates of the no-lag model.

Cost: Same Day Shipping

As discussed in the body of the paper, the lack of a convenience effect could be because our data period does not include significant changes to shipping times. However, anecdotal evidence and Figure 2 indicate that Amazon started to place FCs in locations close to large markets in anticipation of introducing same-day delivery. In ignoring this effect of FC expansion, we may be incorrectly attributing an expected boost in demand due to same-day shipping to cost savings from being close consumers. Intuitively, this would lead to an overestimate of the shipping cost parameter.

In order to address this issue, we add a preference for same-day shipping after 2014 for counties which have access to this service and then re-estimate the supply model. In order to determine the counties which would have same-day delivery, we use a tool on Amazon.com to find the markets which have same day shipping as as of the Fall of 2015.[46] This information allows us to construct a measure of the radius from a FC in which local customers have access to same day delivery, which is about 50 miles. We then add the term $\gamma^{same}\mathbf{1}_{ijt}^{same}$ to ζ_{ijt}, where $\mathbf{1}_{ijt}^{same}$ is equal to one for counties within 50 miles of a FC in year $t \geq 2014$. Because we cannot estimate γ^{same} directly, we choose reasonable values for it and check if/how this changes the results. We choose $\gamma^{same} \in \{0.10, 0.20\}$, which increase the expenditures on Amazon due to same-day shipping by approximately 10 and 20%, respectively. These values are in line with the results in Table 6, which suggest that the effect of 'Local-Express Delivery', the early version of Same-Day delivery, is 0.186.

Results from these robustness estimates under the no lag assumption are in rows four and five of Table 17. We focus on the first column, where the instruments are the indicator variables based on the first stage regression. The first three rows display the baseline estimates, where the direct comparison is to the second row. The downward shift in the confidence set is due to the fact that FC entry now results in both a demand and a supply side effect. That is, the loss in revenue due to the tax effect and an increase in fixed costs is offset by both a reduction in shipping cost and an increase in demand due to same-day shipping.

However, the fact that the confidence sets do not include zero provides reassurance that even if there exists a same day shipping effect, there is still significant cost savings associated with expansion. Therefore, it is likely that economies of density are a key component in determining FC locations.

Cost: Logistics Model

In the primary estimates, we assume that the closest fulfillment center to county i handles all of the shipping to that county. This may be incorrect for a few reasons. First, the lowest cost FC may not be the same as the closest FC. For example, it may be cheaper to ship items to Minneapolis from the FC in Indiana than from the FC in Wisconsin. Presumably Amazon would take this into account when making their distribution decisions. Second, there may be capacity constraints implying that not all shipments can come from a small local FC. Third, it may be the case that the universe of products is not available across all FCs. This last concern is somewhat mitigated by the fact that FCs of different types are often clustered.

In order to address this, we estimate the supply side of the model considering that the shipment of products to county i may be distributed across the FCs in year t and assume that this distribution

[46] `http://www.amazon.com/b?node=8729023011`.

is a function of both the cost of shipping between the FC and county i and the capacity of the FC relative to the total capacity. The relative capacity captures the fact that a single FC may not be able to satisfy all the local demand and the fact that not all types of products will be available at all FCs. The distribution function determines the percentage of expenditures by county i in year t that are shipped from fulfillment center f:

$$\omega_{ift} = \frac{exp(-\beta_1 cost_{if} + \beta_2 cap_f)}{\sum_{f\in F_t} exp(-\beta_1 cost_{if} + \beta_2 cap_f)}$$

where $cost_{if}$ is the estimated log of the average cost of shipping an item between county i and FC f, cap_f is the capacity of FC f in millions of square feet, and F_t is the set of FCs that are open in year t. The estimate average shipping cost is formed using data and model...

Given this distribution, equation 5 becomes:

$$\Pi(a;\theta) = \sum_{t=2006}^{\infty} \beta^{t-2006} \sum_i \mu R_{it}(a_t) - \left(\sum_{f\in F_t} \theta d_{ift}(a_t) \times \omega_{ift} \times R_{it}(a_t) \right) - F_{it}a_{ti} - S_t(a_{it} - a_{it-1}) \tag{15}$$

where d_{ift} is the distance between FC f and county i. Here, the total shipping distance is equivalent to a weighted average of the distance from the FCs in the network, where the weights are based on the percentage of goods coming from each FC.

Again, because we cannot estimate this distribution directly, we choose reasonable values of β_1 and β_2 and check if/how this affects the results. In practice, we choose three different values for β_1 and two different values for β_2, resulting in six additional specifications. Table 16 displays the summary statistics for the shipping distances and the ωs for the closes FC under each combination of the parameters. The values for the base model (closest FC) are in the first row for comparison. A low value of β_1 puts less emphasis on the shipping cost from each FC, which results in shipping being more uniformly distributed around the country. This increases the average shipping distance and the amount of goods being shipped from the closest FC. As β_1 increases, deliveries get more concentrated to the lowest cost FC, which lowers the shipping distance and increases the amount of good being shipped from the closest FC. Note that when $\beta_1 = 200$ and $\beta_2 = 0$, the average distance is close to that of the base model and the majority of shipments are coming from the closest FC.

Increasing the value of β_2 increases the amount of goods that are going to come from the biggest FCs. Note that a value of $\beta_2 = 20$ and $\beta_1 = 0$ results in the distribution of goods shipped from FC f is approximately equal to FC f's capacity relative to total capacity. Therefore, we can think of the model when $\beta_2 = 0$ as a model where capacity does not play a role and the model when $\beta_2 = 20$ as when the amount of goods shipped from FC f is proportional to its capacity. We can see that increasing β_2 increase the average shipping distance and decreases the amount of goods coming from the closest FC.

Overall, under any specification, it is clear that expansion of the network leads to a decrease in the shipping distance to the consumers, implying that there is still a tradeoff between the tax and shipping effects of entering a new state. However, it is not clear how these alternative assumptions will affect the estimates of the supply side model: it depends on the relationship between the change in shipping distance, fixed costs, and revenue as a result of specific FC entry, which may vary across these different models.

Results from these models appear in the last 6 rows of 17. Again focusing on column 1, we see that the interval shifts downwards, which is likely because there is not as strong as a connection between the tax implications and the cost savings of FC entry. However, because the results do not vary significantly from the base model, we believe that the simple model of logistics is a good approximation of the true data generating process.

C.1 Cost: Continuous Instruments

In our base model, the instruments are defined base on whether or not perturbations fall into experiment 1 or experiment 2, implying that there are a number of perturbations that receive 0 weight and, thus, are not used. Therefore, in the second column of Table 17 we present results for our primary models along with the robustness models when including the continuous measure of distance as an instrument. In these specifications, all 1,575 perturbations have weight in the objective function. The results indicated by a star are ones in which the method of Andrews and Soares (2010) results in an empty set. Therefore, we display the value of θ which minimizes the objective function.

The confidence sets shift down because there is now positive weight on perturbations which do not fall into experiments 1 or 2. However, even under the most conservative of estimates, there is still a positive level of cost savings, although they are quite small.

Table 15: Cost Estimates Under State Lag

Instruments	95% Confidence Set	# Perturbations (Group 1)	# Perturbations (Group 2)
	Base: No First Stage		
$\tilde{z}_a$	[5.39E-05,1.67E-04]	313	505
$\tilde{z}_a, \tilde{z}_a \times \tilde{x}_a^+$	[6.11E-05,1.31E-04]	313	505
	Base: First Stage		
$\hat{z}_a$	[4.28E-05,2.63E-04]	272	588
$\hat{z}_a, \hat{z}_a \times \hat{x}_a^+$	[3.89E-05,1.61E-04]	272	588
	Lowest Cost FC: First Stage		
$\hat{z}_a$	[5.46E-05,9.36E-05]	220	576
$\hat{z}_a, \hat{z}_a \times \hat{x}_a^+$	[5.76E-05,8.43E-05]	220	576

Notes: Confidence sets computed using the methods of Andrews and Soares (2010). The instruments indicated by $\hat{z}$ are determined using a first-stage regression. The last two columns indicate the number of perturbations that belong to each experiment. The third panel assumes that Amazon uses the lowest cost FC, rather than the closest, to serve the customers in a county.

Table 16: Alternative Logistics Model

$\{\beta_1, \beta_2\}$	2006	2010	2014	2018
	[Mean Shipping Distance (StDev), Mean ω for Closest FC]			
Closest	[315 (224), 1.00]	[287 (221), 1.00]	[240 (225), 1.00]	[200 (207), 1.00]
$\{10, 0\}$	[366 (231), 0.06]	[357 (228), 0.05]	[352 (233), 0.03]	[338 (219), 0.02]
$\{10, 20\}$	[443 (235), 0.04]	[457 (245), 0.03]	[436 (304), 0.02]	[405 (268), 0.01]
$\{20, 0\}$	[335 (228), 0.07]	[318 (225), 0.07]	[290 (227), 0.05]	[271 (214), 0.04]
$\{20, 20\}$	[397 (234), 0.05]	[400 (238), 0.05]	[347 (259), 0.03]	[311 (232), 0.02]
$\{200, 0\}$	[320 (226), 0.84]	[295 (221), 0.84]	[252 (224), 0.74]	[241 (222), 0.52]
$\{200, 20\}$	[324 (226), 0.78]	[295 (220), 0.84]	[255 (222), 0.70]	[237 (216), 0.42]

Notes: Displayed are the average shipping distance across US counties under different assumptions about Amazon's logistics model. Also displayed are the percentage of volume coming from the closest fulfillment center for each model.

Table 17: Robustness (Cost)

Model	Confidence Set	
	$\hat{z}_a$	$\hat{z}_a, \hat{x}_a^+$
Base: No First Stage	[3.43E-05,1.04E-04]	[2.34E-05,7.05E-05]
Base: First Stage	[5.45E-05,2.26E-04]	[4.65E-05,1.02E-04]
Lowest Cost FC	[3.90E-05,1.20E-04]	[4.29E-05,4.29E-05]*
Sameday ($\beta = 0.10$)	[1.06E-05,1.38E-04]	[1.35E-08,6.15E-05]
Sameday ($\beta = 0.20$)	[2.28E-05,1.28E-04]	[1.36E-08,5.42E-05]
Distribution ($\{\beta_1 = 10, \beta_2 = 0\}$)	[5.73E-05,9.99E-05]	[4.46E-05,4.46E-05]*
Distribution ($\{\beta_1 = 10, \beta_2 = 20\}$)	[2.50E-05,1.47E-04]	[2.15E-05,1.87E-04]
Distribution ($\{\beta_1 = 20, \beta_2 = 0\}$)	[4.08E-05,1.52E-04]	[3.49E-05,3.49E-05]*
Distribution ($\{\beta_1 = 20, \beta_2 = 20\}$)	[3.28E-05,9.53E-05]	[2.96E-05,7.39E-05]
Distribution ($\{\beta_1 = 200, \beta_2 = 0\}$)	[3.92E-05,1.14E-04]	[4.25E-05,4.25E-05]*
Distribution ($\{\beta_1 = 200, \beta_2 = 20\}$)	[4.35E-05,1.08E-04]	[4.53E-05,4.53E-05]*

Notes: Displayed are the estimated 95% confidence sets under different supply and demand side assumptions. The first column uses the indicator variables as instruments, while the second column adds the continuous instruments. A * indicates specifications in which the methods of Andrews and Soares (2010) resulted in an empty set, so the value of θ which minimizes the objective function is displayed.